

Computational Comonads and Intensional Semantics

*Stephen Brookes** *Shai Geva**

Abstract

We explore some foundational issues in the development of a theory of intensional semantics, in which program denotations may convey information about computation strategy in addition to the usual extensional information. Beginning with an “extensional” category \mathcal{C} , whose morphisms we can think of as functions of some kind, we model a notion of computation using a comonad with certain extra structure and we regard the Kleisli category of the comonad as an intensional category. An intensional morphism, or algorithm, can be thought of as a function from computations to values, or as a function from values to values equipped with a computation strategy. Under certain rather general assumptions the underlying category \mathcal{C} can be recovered from the Kleisli category by taking a quotient, derived from a congruence relation that we call extensional equivalence. We then focus on the case where the underlying category is cartesian closed. Under further assumptions the Kleisli category satisfies a weak form of cartesian closure: application morphisms exist, currying and uncurrying of morphisms make sense, and the diagram for exponentiation commutes up to extensional equivalence. When the underlying category is an ordered category we identify conditions under which the exponentiation diagram commutes up to an inequality. We illustrate these ideas and results by introducing some notions of computation on domains and by discussing the properties of the corresponding categories of algorithms on domains.

*School of Computer Science, Carnegie Mellon University, Pittsburgh, Pa 15213, USA.
This research was supported in part by National Science Foundation grant CCR-9006064
and in part by DARPA/NSF grant CCR-8906483.

1 Introduction

Most existing denotational semantic treatments of programming languages are extensional, in that they abstract away from computational details and ascribe essentially extensional meanings to programs. For instance, in the standard denotational treatment of imperative while-programs the meaning of a program is taken to be a partial function from states to states; and in the standard denotational model of the simply typed λ -calculus, the meaning of a term of function type is taken to be a continuous function. Extensional models are appropriate if one wants to reason only about extensional properties of programs, such as partial correctness of while-programs. However, such models give no insight into questions concerning essentially intensional aspects of program behavior, such as efficiency or complexity. For instance, in a typical extensional model all sorting programs denote the same function, regardless of their computation strategy, and therefore regardless of their worst- or average-case behavior. It is desirable to have a semantic model in which sensible comparisons can be made between programs with the same extensional behavior, on the basis of their computation strategy.

We emphasize that we regard intensionality as a relative term; given a programming language we might wish to provide an extensional semantics and also an intensional semantics that contains more computational information and is thus at a lower level of abstraction. We would like to be able to extract extensional meanings from intensional meanings, and to show that the intensional semantics “fits properly” on top of the extensional semantics. Suppose that we have an extensional semantics provided in a category whose objects represent sets of data values and whose morphisms are functions of some kind; and that we have an intensional semantics in a category with the same objects but with morphisms that we regard as algorithms, which correspond to functions equipped with a computation strategy. It seems reasonable that we should be able to define an equivalence relation on algorithms (in the same hom-set) that identifies all pairs of algorithms with the same “extensional part”; that composition of algorithms should respect this equivalence; and that quotienting the algorithms from A to B by this equivalence relation should yield precisely the extensional morphisms.

In this paper we set out a basis for a category-theoretic approach to intensional semantics, motivated by the following intuition. If the extensional meaning of a program may be modelled as some kind of function from data values to data values, then we can obtain an intensional semantics by introducing a notion of computation and defining an intensional meaning to be a function from computations to values. This accords with an intuitive operational semantics for programs in which the execution of a program proceeds lazily in a demand-driven, coroutine-like manner [10]: the program responds

to requests for output (say, from a user) by performing input computation until it has sufficient information to supply an output value. We formalize what we mean by a notion of computation in abstract terms as follows. Suppose that extensional meanings are given in some category \mathcal{C} . Then, for each object A , we specify an object TA of computations over A and we specify how to lift a morphism f from A to B into a morphism Tf from TA to TB ; we require that T be a functor on \mathcal{C} . We specify, for each object A , a morphism $\epsilon_A : TA \rightarrow A$ from computations to values and a morphism $\delta_A : TA \rightarrow T^2A$ that maps a computation over A to a computation over TA . Intuitively, ϵ maps a computation to the value it computes, and δ shows how a computation may itself be computed. We require that (T, ϵ, δ) be a comonad over \mathcal{C} . Then we regard the Kleisli category of this comonad as an intensional category; it has the same objects as \mathcal{C} , and an intensional morphism from A to B is just a morphism in \mathcal{C} from TA to B .

We say that a comonad is computational if there is a natural way to convert a data value into a degenerate computation returning that value. This enables us to extract from an algorithm a function from values to values, and we obtain an extensional equivalence relation on algorithms by identifying all pairs of algorithms that determine the same function. We show that if the comonad is computational then the Kleisli category collapses onto \mathcal{C} under extensional equivalence. This means that an intensional semantics based on a computational comonad can be viewed as a conservative extension of a corresponding extensional semantics based on the underlying category. We may therefore use this framework for reasoning at different levels of abstraction, while being assured that the addition of intensional information does not interfere with purely extensional aspects of program behavior.

We then show that, assuming certain further conditions concerning products, the Kleisli category satisfies an intensional analogue of the cartesian closedness property. This generalizes from the known result that if the underlying category is cartesian closed and the (functor part of the) comonad preserves products then the Kleisli category is also cartesian closed. When the underlying category \mathcal{C} is an ordered category, we identify conditions under which the Kleisli construction preserves certain lax forms of cartesian closedness.

Throughout the paper we motivate our definitions and results by means of notions of computation on domains. We focus primarily on three forms of computation at differing levels of abstraction. At the end of the paper we discuss briefly some further examples that indicate the broader applicability of our ideas. We also mention the relationship between our work and the sequential algorithms model of Berry and Curien [6].

We assume familiarity with elementary category theory and domain theory.

We refer the reader to [11] and [1] for categorical background and to [8] for the relevant domain theory.

2 Computations, Comonads and Algorithms

Let \mathcal{C} be a category that we regard as providing an extensional framework. We wish to encapsulate in abstract terms what a notion of computation over \mathcal{C} is, and to build an “intensional” category whose morphisms can be thought of as extensional morphisms equipped with additional computational information. We model a notion of computation over \mathcal{C} using a comonad over \mathcal{C} , the functor part of which maps an object A to an object TA representing computations over A . The two other components of the comonad describe how to extract a value from a computation, and how a computation is built up from its sub-computations. We then take an intensional morphism from A to B to be an extensional morphism from TA to B , essentially a morphism from *input computations* over A to *output values* in B . This leads us to use for our intensional category the Kleisli category \mathcal{C}_T [11], which has the same objects as \mathcal{C} and in which the morphisms from A to B are exactly the \mathcal{C} -morphisms from TA to B . Typically \mathcal{C} is a category in which morphisms are functions of some kind, and we will refer to intensional morphisms in \mathcal{C}_T as *algorithms* to emphasize their computational content. In case we need to compare Kleisli categories for different comonads over the same underlying category we will use the term T -algorithm, indicating the comonad explicitly.

2.1 Comonads and the Kleisli category

Definition 2.1 A *comonad* over a category \mathcal{C} is a triple (T, ϵ, δ) where $T : \mathcal{C} \rightarrow \mathcal{C}$ is a functor, $\epsilon : T \rightarrow I_{\mathcal{C}}$ is a natural transformation from T to the identity functor, and $\delta : T \rightarrow T^2$ is a natural transformation from T to T^2 , such that the following associativity and identity conditions hold, for every object A :

$$\begin{aligned} T(\delta_A) \circ \delta_A &= \delta_{TA} \circ \delta_A \\ \epsilon_{TA} \circ \delta_A &= T(\epsilon_A) \circ \delta_A = \text{id}_{TA}. \end{aligned}$$

Figures 1 and 2 express these requirements in diagrammatic form. •

Definition 2.2 Given a comonad (T, ϵ, δ) over \mathcal{C} , the *Kleisli category* \mathcal{C}_T is defined by:

- The objects of \mathcal{C}_T are the objects of \mathcal{C} .

BROOKES AND GEVA : COMPUTATIONAL COMONADS




Figure 1. Naturality of ϵ and δ in a comonad: these diagrams commute, for all $A, B, f : A \rightarrow^c B$.




Figure 2. Identity and associativity laws of a comonad: these diagrams commute, for all A .

- The morphisms from A to B in \mathcal{C}_T are the morphisms from TA to B in \mathcal{C} .
- The identity morphism $\widehat{\text{id}}_A$ on A in \mathcal{C}_T is $\epsilon_A : TA \rightarrow^{\mathcal{C}} A$.
- The composition in \mathcal{C}_T of $a : A \rightarrow^{\mathcal{C}_T} B$ and $a' : B \rightarrow^{\mathcal{C}_T} C$, denoted $a' \bar{\circ} a$, is the composition in \mathcal{C} of $\delta_A : TA \rightarrow^{\mathcal{C}} T^2A$, $Ta : T^2A \rightarrow^{\mathcal{C}} TB$ and $a' : TB \rightarrow^{\mathcal{C}} C$, i.e.,

$$a' \bar{\circ} a = a' \circ Ta \circ \delta_A.$$

The associativity and identity laws of the comonad ensure that \mathcal{C}_T is a category [11]. •

This use of morphisms from TA to B to model algorithms from A to B fits well with an intuitive operational semantics based on the coroutine mechanism [10]. A program responds to requests for output by requiring some computation on its input (typically, to evaluate some portion of the input) until it has enough information to determine what output value to produce. Execution is lazy, in that computation is demand-driven. The operational behavior of algorithm composition can be described as follows. Let $a : A \rightarrow^{\mathcal{C}_T} B$ and $a' : B \rightarrow^{\mathcal{C}_T} C$. Then $a' \bar{\circ} a$ responds to a request for output (in C) by performing an input computation t over A , transforming this into a computation t' over B by applying δ and then Ta to t , and supplying t' as argument to a' . For further details concerning operational semantics we refer to [5].

3 Notions of computation on domains

Our main examples will be based on a category of domains and continuous functions. To avoid repetition and to be precise, let us remark that by a domain we mean a directed-complete, bounded-complete, algebraic partial order with a least element. That is, a domain (D, \sqsubseteq) is a set D equipped with a partial order \sqsubseteq satisfying the following conditions:

- D has a least element, denoted \perp_D .
- Every non-empty directed subset $X \subseteq D$ has a least upper bound $\bigsqcup X$.
- Every non-empty bounded subset $X \subseteq D$ has a least upper bound $\bigsqcup X$.
- Every element of D is the least upper bound of its (directed set of) finite approximations.

A set X is directed iff for all $x, y \in X$ there is a $z \in X$ such that $x \sqsubseteq z$ and $y \sqsubseteq z$. A set X is bounded (or consistent) iff there is a $z \in D$ such that $x \sqsubseteq z$ for all $x \in X$. An element $e \in D$ is finite (or isolated) iff, for all directed subsets $X \subseteq D$, if $e \sqsubseteq \bigsqcup X$ then $e \sqsubseteq x$ for some $x \in X$.

We remark that none of our example comonads really requires that the underlying domain be algebraic, and nor does the presence of a least element play any prominent rôle (except, of course, in justifying the existence of least fixed points). We could just as well work in the category of directed-complete, bounded complete partial orders and continuous functions. Nevertheless, the property of algebraicity is very natural in the computational setting and all of our example functors on domains preserve algebraicity. At the end of the paper we will discuss further examples based on different categories and different types of domain.

3.1 Increasing paths

The first notion of computation that we introduce models in abstract terms a sequence of time steps in which some incremental evaluation is being performed. For example, a program with two inputs may need to evaluate one or more of its input arguments and it may attempt to perform evaluations in parallel or in sequence; moreover, it may only require partial information about its arguments, as is typically the case, say, when an argument is a function and the program needs to apply that argument to an already known parameter. One natural way to formalize this form of computation is as an increasing sequence of values drawn from a domain, whose partial order reflects the amount of information inherent in a value.

We define the comonad T_1 of “increasing paths” as follows¹.

- For a given domain (D, \sqsubseteq) , let T_1D be the set of infinite increasing sequences over D , ordered componentwise. Thus, the elements of T_1D have form $\langle d_n \rangle_{n=0}^\infty$, where for each $n \geq 0$, $d_n \sqsubseteq d_{n+1}$; and we define $\langle d_n \rangle_{n=0}^\infty \sqsubseteq_{T_1D} \langle d'_n \rangle_{n=0}^\infty$ iff for all $n \geq 0$, $d_n \sqsubseteq_D d'_n$.
- For a continuous function $f : D \rightarrow D'$, let $T_1f : T_1D \rightarrow T_1D'$ be the function that applies f componentwise. That is, $(T_1f)\langle d_n \rangle_{n=0}^\infty = \langle fd_n \rangle_{n=0}^\infty$.
- For $t \in T_1D$ let $\epsilon_D t$ be the least upper bound of t . That is, $\epsilon\langle d_n \rangle_{n=0}^\infty = \bigsqcup_{n=0}^\infty d_n$.

¹This comonad, adapted for Scott domains, was first introduced in [4].




Figure 3. The domain $\mathbf{Two} \times \mathbf{Two}$.

- For $t \in T_1D$ let $\delta_D t$ be the sequence of prefixes of t , each represented as an infinite sequence. That is, if $t = \langle d_n \rangle_{n=0}^\infty$, then for each $n \geq 0$, $(\delta t)_n = d_0 \dots d_n d_n^\omega$.

Intuitively, a computation may be viewed as a (time-indexed) sequence of increments in the amount of information known about a data value, and the value “computed” by such a computation is its least upper bound; each computation is itself built up progressively from its prefixes. Equivalently, we can regard a computation over D as a continuous function from the domain \mathbf{VNat} of “vertical” natural numbers (together with limit point ω) to D . Our ordering on computations then corresponds exactly to the pointwise ordering on such functions.

The least element of T_1D is the sequence \perp^ω . The finite (isolated) elements of T_1D are just the eventually constant sequences all of whose elements are finite in D . It is easy to verify that T_1 maps domains to domains and is indeed a functor. The comonad laws hold: naturality of ϵ corresponds to continuity of f ; naturality of δ states that the operation of applying a function componentwise to a sequence “commutes” with taking prefixes; every computation is the least upper bound of its prefixes; and every prefix of a prefix of t is also a prefix of t .

For illustration, let \mathbf{Two} be the domain $\{\perp, \top\}$. The domain $\mathbf{Two} \times \mathbf{Two}$ is shown in Figure 3, and Figure 4 shows the six continuous functions from $\mathbf{Two} \times \mathbf{Two}$ to \mathbf{Two} , ordered pointwise. We give these functions mnemonic names: \perp and \top are constant functions; l is strict in its left argument; r is strict in its right argument; b is strict in both arguments; poll returns \top if either of its two arguments is \top , so that poll is not strict in either argument. Each function is depicted by a Hasse diagram corresponding to Figure 3, in which the nodes are shaded to indicate their image under the function being described: \circ corresponds to \perp , \bullet to \top .

Figure 5 shows part of $T_1(\mathbf{Two} \times \mathbf{Two})$. Figure 6 shows some of the T_1 -algorithms from $\mathbf{Two} \times \mathbf{Two}$ to \mathbf{Two} , ordered pointwise. The notation for describing algorithms is based on Figure 5, again with \circ representing \perp and \bullet representing \top . In each case the intended algorithm is the least continuous

BROOKES AND GEVA : COMPUTATIONAL COMONADS

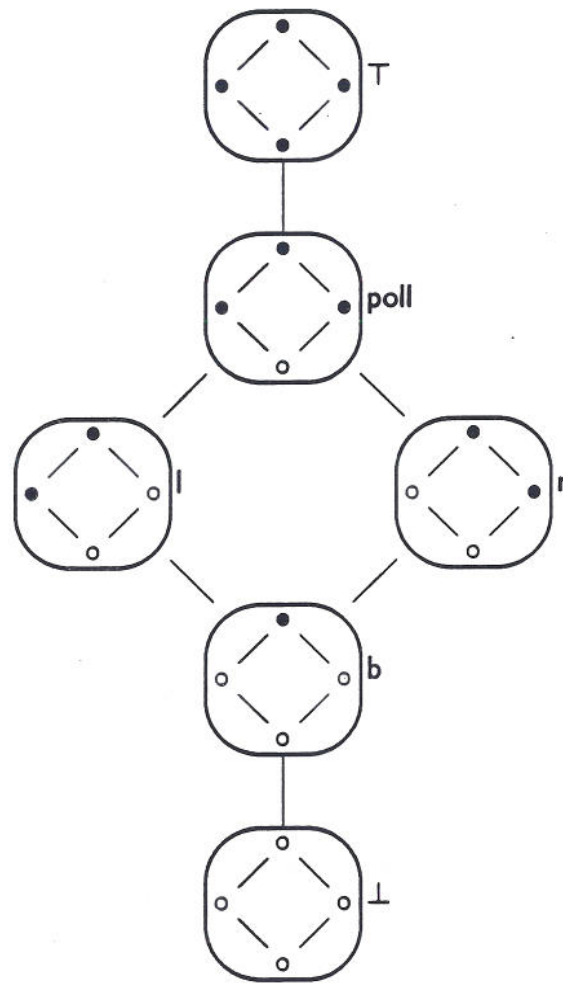
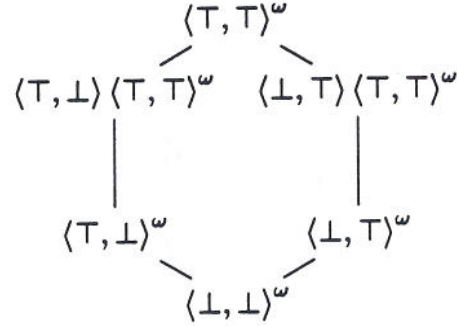


Figure 4. Continuous functions from $\mathbf{Two} \times \mathbf{Two}$ to \mathbf{Two} , ordered pointwise.


 Figure 5. Some paths in $T_1(\text{Two} \times \text{Two})$.

function on paths consistent with this description. The nomenclature is intended to indicate (as yet only informally) the function computed by each algorithm and what computation strategy the algorithm uses. For instance, the algorithms pb , lb , rb and db all compute the function b ; pb computes both arguments in parallel immediately, lb computes left-first and then right, rb computes right-first and then left, and db computes both arguments in either order. Since the diagram includes only one algorithm for poll , for \perp , and for T , in these cases we use the same name for the algorithm as for the function.

Since $T_1(\text{Two} \times \text{Two})$ includes paths with repeated steps, we can also make distinctions between algorithms which differ not in the order in which they evaluate their arguments, but in the amount of time they are prepared to wait for each successive increment to be achieved. For instance, for the function b there are algorithms pb_n , lb_n , rb_n and db_n for each $n \geq 0$, characterized as the least functions on paths such that

$$\begin{aligned} \text{pb}_n(\langle \perp, \perp \rangle^n \langle T, T \rangle^\omega) &= T \\ \text{lb}_n(\langle \perp, \perp \rangle^n \langle T, \perp \rangle \langle T, \perp \rangle^n \langle T, T \rangle^\omega) &= T \\ \text{rb}_n(\langle \perp, \perp \rangle^n \langle \perp, T \rangle \langle \perp, T \rangle^n \langle T, T \rangle^\omega) &= T \\ \text{db}_n(\langle \perp, \perp \rangle^n \langle T, \perp \rangle \langle T, \perp \rangle^n \langle T, T \rangle^\omega) &= T \\ \text{db}_n(\langle \perp, \perp \rangle^n \langle \perp, T \rangle \langle \perp, T \rangle^n \langle T, T \rangle^\omega) &= T. \end{aligned}$$

Informally, pb_n is the algorithm which needs to evaluate both arguments and returns T provided each evaluation succeeds (with result T) in at most n time steps. Similarly, lb_n evaluates both arguments and returns T provided evaluation of the left argument succeeds in at most n time steps and evaluation of the right argument succeeds in at most $2n + 1$ time steps.

The following relationships hold, for all $n \geq 0$:

$$\begin{aligned} \text{pb}_n &\subseteq \text{lb}_n \subseteq \text{pb}_{2n+1} \\ \text{pb}_n &\subseteq \text{rb}_n \subseteq \text{pb}_{2n+1} \\ \text{db}_n &= \text{lb}_n \cup \text{rb}_n. \end{aligned}$$

BROOKES AND GEVA : COMPUTATIONAL COMONADS




Figure 6. Some T_1 -algorithms from $\mathbf{Two} \times \mathbf{Two}$ to \mathbf{Two} , ordered pointwise.

Moreover, $pb_n \sqsubseteq pb_{n+1}$, $lb_n \sqsubseteq lb_{n+1}$, $r_n \sqsubseteq r_{n+1}$, and $db_n \sqsubseteq db_{n+1}$ for all $n \geq 0$. Each of these sequences of algorithms has the same least upper bound, characterized as the algorithm b_* that maps every path with lub $\langle \top, \top \rangle$ to \top . Of course, in Figure 6, pb is just pb_0 , and so on.

3.2 Strictly increasing paths

In the comonad T_1 a computation has a built-in measure of the number of time steps it takes between successive proper increments. We obtain a more abstract notion of computation by retaining only the increments themselves, so that we may still make distinctions on the basis of the order of evaluation of arguments. To do this we model a computation as a “strictly increasing path”. We define the strictly increasing path comonad T_2 as follows.

- Let T_2D be the set of strictly increasing finite and infinite sequences over D . For convenience, we represent a finite sequence as an eventually constant infinite sequence. That is, the elements of T_2D are either of form $\langle d_n \rangle_{n=0}^\infty$, with $d_n \sqsubset_D d_{n+1}$ for all $n \geq 0$; or of form $d_0 \dots d_{N-1} d_N^\omega$, where $N \geq 0$ and $d_n \sqsubset_D d_{n+1}$ for $0 \leq n < N$. Let \sqsubseteq_{T_2D} be the least partial order on T_2D such that


$$d_0 \dots d_{N-1} d_N^\omega \sqsubseteq_{T_2D} d_0 \dots d_{N-1} t \quad \text{if } t \in T_2D \text{ \& } d_N \sqsubseteq_D t_0.$$

This ordering is based on the prefix ordering on sequences, but adjusted to take appropriate account of the underlying order on data values. The order \sqsubseteq_{T_2D} is actually the stable ordering [2] on T_2D , when we regard the elements of T_2D as (strictly increasing, possibly eventually constant) stable functions from $\mathbb{V}\text{Nat}$ to D . Note that every continuous function from $\mathbb{V}\text{Nat}$ to D is also stable.

- For a continuous function $f : D \rightarrow D'$ we define T_2f to be the function which applies f componentwise and suppresses any repetitions (except for constant suffixes). That is, T_2f is the least continuous function such that for all $d \in D$, for all $d_0, d_1 \in D$ such that $d_0 \sqsubset d_1$, and for all $t \in T_2D$,

$$\begin{aligned} T_2f(d^\omega) &= (fd)^\omega \\ T_2f(d_0d_1t) &= (fd_0)(T_2f(d_1t)) && \text{if } fd_0 \neq fd_1 \\ &= T_2f(d_1t) && \text{if } fd_0 = fd_1. \end{aligned}$$

- For all $t \in T_2D$, let $\epsilon_D t$ be the lub of t . That is, $\epsilon \langle d_n \rangle_{n=0}^\infty = \bigsqcup_{n=0}^\infty d_n$.


 Figure 7. Part of $T_2(\text{Two} \times \text{Two})$.

- For all $t \in T_2D$, let $\delta_D t$ be the sequence of prefixes of t . Again, if $t = \langle d_n \rangle_{n=0}^\infty$ then for each $n \geq 0$, $(\delta t)_n = d_0 \dots d_n d_n^\omega$. Note that if t is strictly increasing, so is δt .

The least element of T_2D is the sequence \perp^ω . The lub of a directed (or bounded) set of strictly increasing paths is again a strictly increasing path². The finite elements of T_2D are the eventually constant sequences of form $\langle d_0 \dots d_{N-1} \rangle d_N^\omega$, where $N \geq 0$ and d_N is a finite element of D . Every element of T_2D is the lub of its finite approximations. Thus, T_2 maps domains to domains. Functoriality is easily checked.

Although the order is not pointwise, it is still true that every $t \in T_2D$ is the lub of its prefixes. The comonad laws hold for (T_2, ϵ, δ) .

Figure 7 shows some of the paths in $T_2(\text{Two} \times \text{Two})$. Figure 8 shows some of the T_2 -algorithms from $\text{Two} \times \text{Two}$ to Two , ordered pointwise, using a notation based on Figure 7. Again the nomenclature is chosen to indicate the function computed and the computation strategy. Each of the T_1 -algorithms of Figure 6 has a corresponding T_2 -algorithm, for which we use the same name; but because of the different ordering on paths, there are three additional T_2 -algorithms. Note also that since T_2D does not include paths with repeated elements, only pb_0 and pb_1 of the family of pb_n algorithms defined above have corresponding T_2 -algorithms.

3.3 Timed data

A simple notion of computation over domains is obtained by regarding a computation as a pair consisting of a data value and a natural number, intuitively representing the amount of time or the cost associated with the calculation

²However, this would not be the case if we ordered T_2D componentwise, since T_2D is not directed-complete under the componentwise ordering.

BROOKES AND GEVA : COMPUTATIONAL COMONADS

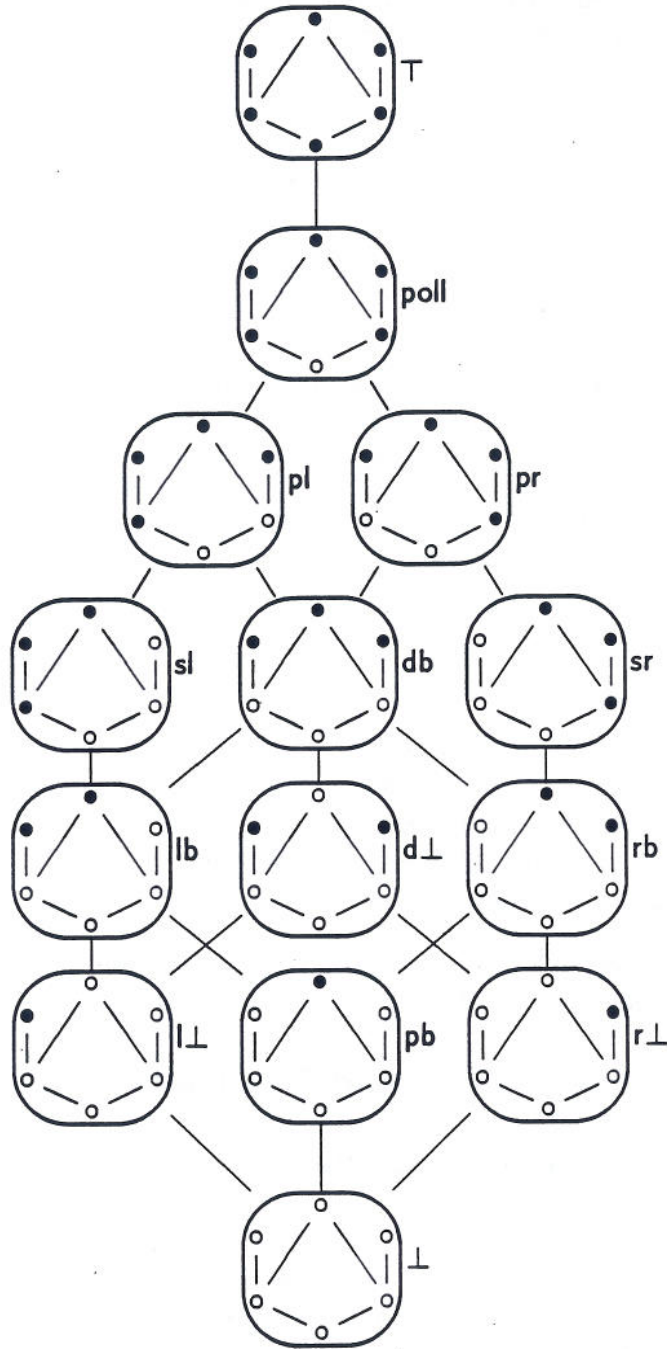


Figure 8. Some T_2 -algorithms from $Two \times Two$ to Two , ordered pointwise.

of the value. With this intuition it seems reasonable to regard one computation $\langle d, n \rangle$ as approximating another $\langle d', n' \rangle$ iff $d \sqsubseteq d'$ and $n' \leq n$; that is, a better computation produces a more precise data value in less time. This suggests the use of the following comonad:

- $T_3D = D \times \text{VNat}^{\text{op}}$, ordered componentwise.
- For $f : D \rightarrow D'$, $(T_3f) \langle d, n \rangle = \langle fd, n \rangle$.
- $\epsilon \langle d, n \rangle = d$.
- $\delta \langle d, n \rangle = \langle \langle d, n \rangle, n \rangle$.

Here VNat^{op} is the domain consisting of the natural numbers together with ω , ordered by the reverse of the usual ordering, so that ω is the least element. The least element of T_3D is $\langle \perp_D, \omega \rangle$.

We may define for each continuous function $f : D \rightarrow D'$ and each $n \in \text{VNat}$ an algorithm f_n from D to D' :

$$\begin{aligned} f_n \langle d, k \rangle &= fd && \text{if } k \leq n \\ &= \perp && \text{otherwise.} \end{aligned}$$

Clearly, whenever $f \sqsubseteq g$ we also get $f_n \sqsubseteq g_n$. Moreover, because of our ordering on T_3D , we get $f_n \sqsubseteq f_{n+1}$ for each $n \geq 0$; and f_ω is simply $\lambda \langle d, n \rangle . fd$. The lub of the f_n is the function $f_* = \lambda \langle d, n \rangle . (n = \omega \rightarrow \perp, fd)$, which is below f_ω . Using this nomenclature, we show some of the T_3 -algorithms from $\text{Two} \times \text{Two}$ to Two in Figure 9.

It is also possible to define a comonad based on the functor $TD = D \times \text{VNat}$, using the usual ordering on the integer component.

4 Relating algorithms and functions

4.1 Computational comonads

We will say that a comonad is *computational* if for each object A , every data value in A can be regarded as a “degenerate” computation in TA , and degenerate computations possess certain simple properties. This permits us to extract from an algorithm a function from values to values, by looking at the algorithm’s effect when applied to degenerate computations. Two algorithms are called extensionally equivalent iff they determine the same function.

More precisely, we define a computational comonad to be a comonad (T, ϵ, δ) together with a natural transformation $\gamma : I_C \rightarrow T$ satisfying some axioms

BROOKES AND GEVA : COMPUTATIONAL COMONADS

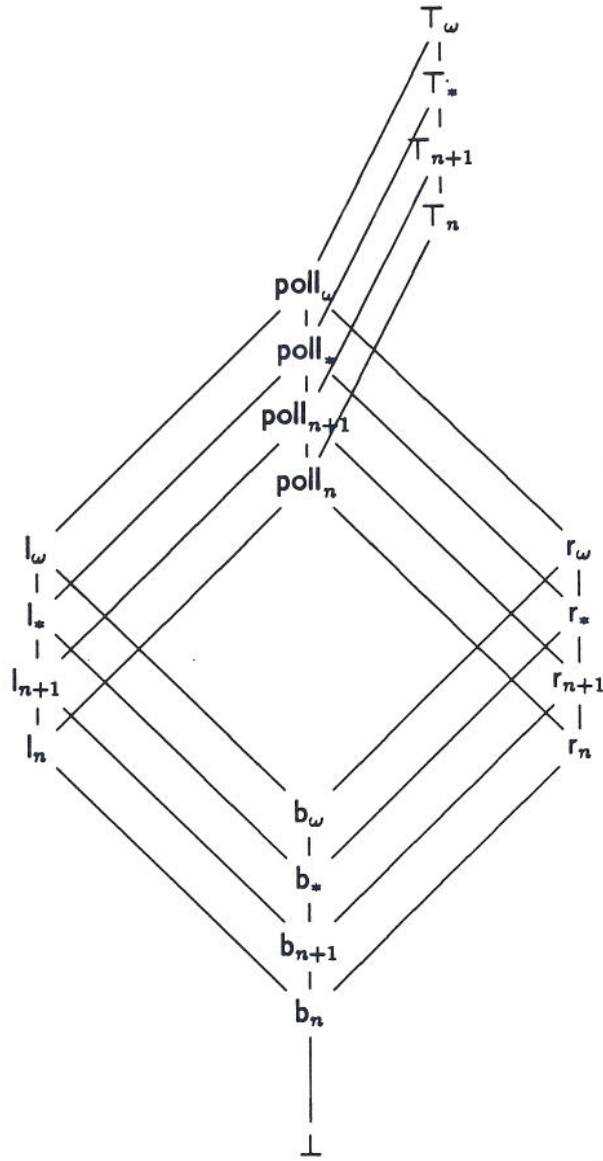


Figure 9. Some T_3 -algorithms from $\mathbf{Two} \times \mathbf{Two}$ to \mathbf{Two} , ordered pointwise.




Figure 10. Properties of a computational comonad: these diagrams commute, for all A, B , and all morphisms $f : A \rightarrow^{\mathcal{C}} B$.

which capture formally what we mean by degeneracy. We then show that this permits us to define an “extensional equivalence” relation on each hom-set in \mathcal{C}_T . Extensional equivalence is preserved by composition, so that we actually have a congruence on \mathcal{C}_T . The underlying category \mathcal{C} may then be recovered from \mathcal{C}_T by taking a quotient.

Definition 4.1 A *computational comonad* over a category \mathcal{C} is a quadruple $(T, \epsilon, \delta, \gamma)$ where (T, ϵ, δ) is a comonad over \mathcal{C} and $\gamma : I_{\mathcal{C}} \rightarrow T$ is a natural transformation such that, for every object A ,

- $\epsilon_A \circ \gamma_A = \text{id}_A$
- $\delta_A \circ \gamma_A = \gamma_{TA} \circ \gamma_A$.

Naturality guarantees that, for every morphism $f : A \rightarrow^{\mathcal{C}} B$,

- $Tf \circ \gamma_A = \gamma_B \circ f$.

Figure 10 shows these properties in diagrammatic form. •

As an immediate corollary of these properties, ϵ_A is epi and γ_A is mono, for every object A .

Proposition 4.2 *If $(T, \epsilon, \delta, \gamma)$ is a computational comonad, then there is a pair of functors (alg, fun) between \mathcal{C} and \mathcal{C}_T with the following definitions and properties:*

BROOKES AND GEVA : COMPUTATIONAL COMONADS

- $\text{alg} : \mathcal{C} \rightarrow \mathcal{C}_T$ is the identity on objects, and $\text{alg}(f) = f \circ \epsilon_A$, for every $f : A \rightarrow^{\mathcal{C}} B$.
- $\text{fun} : \mathcal{C}_T \rightarrow \mathcal{C}$ is the identity on objects, and $\text{fun}(a) = a \circ \gamma_A$, for all $a : A \rightarrow^{\mathcal{C}_T} B$.
- fun induces an equivalence relation $=^e$ on \mathcal{C}_T , given by

$$a_1 =^e a_2 \iff \text{fun}(a_1) = \text{fun}(a_2).$$

This relation is a congruence; that is, for all $a_1, a_2 : A \rightarrow^{\mathcal{C}_T} B$ and $a'_1, a'_2 : B \rightarrow^{\mathcal{C}_T} C$,

$$a_1 =^e a_2 \ \& \ a'_1 =^e a'_2 \ \Rightarrow \ a'_1 \bar{\circ} a_1 =^e a'_2 \bar{\circ} a_2.$$

- The quotient category of \mathcal{C}_T by $=^e$ is isomorphic to \mathcal{C} .
- $\text{fun} \circ \text{alg} = I_{\mathcal{C}}$. That is, for all $f : A \rightarrow^{\mathcal{C}} B$, $\text{fun}(\text{alg } f) = f$.
- $\text{alg} \circ \text{fun} =^e I_{\mathcal{C}_T}$, in that for all $a : A \rightarrow^{\mathcal{C}_T} B$, $\text{alg}(\text{fun } a) =^e a$.

Proof: Functoriality of alg and fun are straightforward. For instance:

$$\begin{aligned} \text{fun}(a' \bar{\circ} a) &= (a' \bar{\circ} a) \circ \gamma \\ &= (a' \circ T a \circ \delta) \circ \gamma \\ &= a' \circ T a \circ \gamma \circ \gamma \quad \text{since } \delta \circ \gamma = \gamma \circ \gamma \\ &= a' \circ \gamma \circ a \circ \gamma \quad \text{by naturality of } \gamma \\ &= \text{fun}(a') \circ \text{fun}(a). \end{aligned}$$

A similar calculation shows that $=^e$ is a congruence.

The quotient of \mathcal{C}_T by $=^e$ has the same objects as \mathcal{C}_T (and therefore the same objects as \mathcal{C}), and the morphisms in the quotient category from A to B are the $=^e$ -equivalence classes of morphisms from A to B in \mathcal{C}_T . Let us write $[a]$ for the equivalence class of a . Clearly, the map $f \mapsto [\text{alg}(f)]$ is an isomorphism of hom-sets, showing that $\mathcal{C}_T / =^e$ is isomorphic to \mathcal{C} .

The facts that $\text{fun} \circ \text{alg} = I_{\mathcal{C}}$ and $\text{alg} \circ \text{fun} =^e I_{\mathcal{C}_T}$ are elementary consequences of the definitions. \square

We say that $\text{fun}(a)$ is the extensional morphism computed by a . Since $\text{fun}(\text{alg } f) = f$, every extensional morphism f is computed by some (not necessarily unique) intensional morphism.

$$\begin{array}{ccc}
 \mathcal{C}(TA, B) & \xrightarrow{\mathcal{C}(Tf, g)} & \mathcal{C}(TA', B') \\
 \text{fun} \downarrow & & \downarrow \text{fun} \\
 \mathcal{C}(A, B) & \xrightarrow{\mathcal{C}(f, g)} & \mathcal{C}(A', B')
 \end{array}
 \qquad
 \begin{array}{ccc}
 \mathcal{C}(TA, B) & \xrightarrow{\mathcal{C}(Tf, g)} & \mathcal{C}(TA', B') \\
 \text{alg} \uparrow & & \uparrow \text{alg} \\
 \mathcal{C}(A, B) & \xrightarrow{\mathcal{C}(f, g)} & \mathcal{C}(A', B')
 \end{array}$$

Figure 11. Naturality of fun and alg for a computational comonad: these diagrams commute, for all $f : A' \rightarrow^{\mathcal{C}} A$ and $g : B \rightarrow^{\mathcal{C}} B'$.

These results show that every computational comonad can be used to produce an intensional category that yields back the underlying extensional category when we take the extensional quotient. Next we show that fun and alg are natural transformations. Let Set be the category of sets and functions.

Definition 4.3 The two-variable hom-functor $\mathcal{C}(T(-), -)$ from $\mathcal{C}^{\text{op}} \times \mathcal{C}$ to Set takes a pair (A, B) of objects to $\mathcal{C}(TA, B)$ and takes a pair of morphisms (f, g) with $f : A' \rightarrow^{\mathcal{C}} A$ and $g : B \rightarrow^{\mathcal{C}} B'$ to $\mathcal{C}(Tf, g) : \mathcal{C}(TA, B) \rightarrow \mathcal{C}(TA', B')$, where for all $a : TA \rightarrow^{\mathcal{C}} B$,

$$\mathcal{C}(Tf, g)(a) = g \circ a \circ Tf.$$

Similarly, the two-variable hom-functor $\mathcal{C}(-, -)$ takes (A, B) to $\mathcal{C}(A, B)$ and (f, g) to $\mathcal{C}(f, g)$, where for all $h : A \rightarrow^{\mathcal{C}} B$,

$$\mathcal{C}(f, g)(h) = g \circ h \circ f.$$

Proposition 4.4 Let $(T, \epsilon, \delta, \gamma)$ be a computational comonad over a category \mathcal{C} . Then fun and alg , as defined in Proposition 4.2, are natural transformations:

$$\begin{aligned}
 \text{fun} &: \mathcal{C}(T(-), -) \rightarrow \mathcal{C}(-, -) \\
 \text{alg} &: \mathcal{C}(-, -) \rightarrow \mathcal{C}(T(-), -).
 \end{aligned}$$

That is, for all $f : A' \rightarrow^c A$ and $g : B \rightarrow^c B'$, the following identities hold:

$$\begin{aligned} \text{fun} \circ \mathcal{C}(Tf, g) &= \mathcal{C}(f, g) \circ \text{fun} \\ \mathcal{C}(Tf, g) \circ \text{alg} &= \text{alg} \circ \mathcal{C}(f, g). \end{aligned}$$

Figure 11 expresses these properties in diagram form.

Proof: Straightforward, using naturality of γ and ϵ . □

We now show that the three example comonads introduced earlier can be extended to become examples of computational comonads.

4.2 Examples.

1. For the increasing paths comonad T_1 , let $\gamma_D : D \rightarrow T_1 D$ be defined by $\gamma_D d = d^\omega$, for all $d \in D$. Clearly γ_D is continuous, and γ is a natural transformation. Moreover, the computational comonad laws hold, since the lub of d^ω is d , and all prefixes of d^ω are equal to d^ω .

The functor fun maps each algorithm from $\text{Two} \times \text{Two}$ to Two into a function from $\text{Two} \times \text{Two}$ to Two . In particular, $\text{fun}(\text{pb}) = \text{fun}(\text{lb}) = \text{fun}(\text{rb}) = \text{fun}(\text{db}) = \text{b}$. Similarly, $\text{fun}(\text{pb}_n) = \text{b}$ for all $n \geq 0$, and $\text{fun}(\text{b}_*) = \text{b}$. In fact, $\text{b}_* = \text{alg}(\text{b})$.

Figure 12 illustrates the result of taking the extensional quotient of Figure 6. Boxes enclose equivalence classes of algorithms, arcs between boxes represent the quotient ordering, and within each box we retain the pointwise order to ease comparison with Figure 6. As expected, the quotient figure is isomorphic to Figure 4 when we identify each equivalence class with the function computed.

2. For the strictly increasing paths comonad, we again let $\gamma_D : D \rightarrow T_2 D$ be $\gamma_D d = d^\omega$. Again this is a continuous function, and γ is a natural transformation. Again the computational comonad laws hold. Figure 13 shows the quotient of Figure 8 by extensional equivalence. Note that $\text{fun}(d\perp) = \text{fun}(l\perp) = \text{fun}(r\perp) = \text{fun}(\perp) = \perp$.

Again the quotient diagram is isomorphic to Figure 4.

3. For the timed data comonad, we obtain a suitable γ by deciding what cost to associate with a degenerate computation. For each $k \in \text{VNat}$ we may take $\gamma_k d = \langle d, k \rangle$ and obtain a natural transformation satisfying the requirements of a computational comonad. Define fun_k to be the functor whose action on algorithms is given by $\text{fun}_k(a) = a \circ \gamma_k$, and let

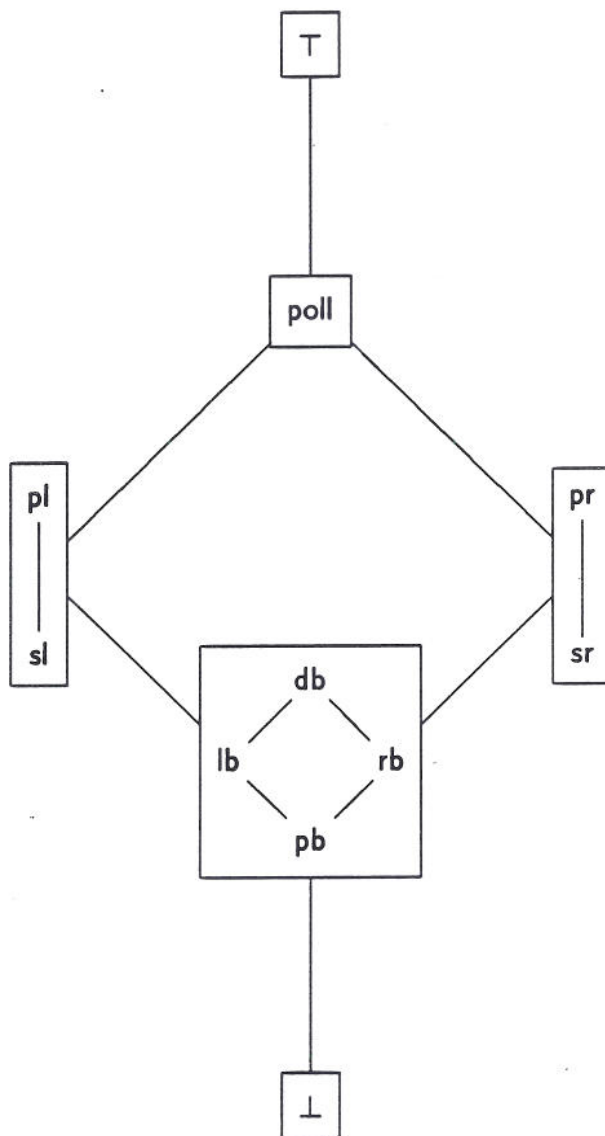


Figure 12. Quotient of Figure 6 by extensional equivalence.

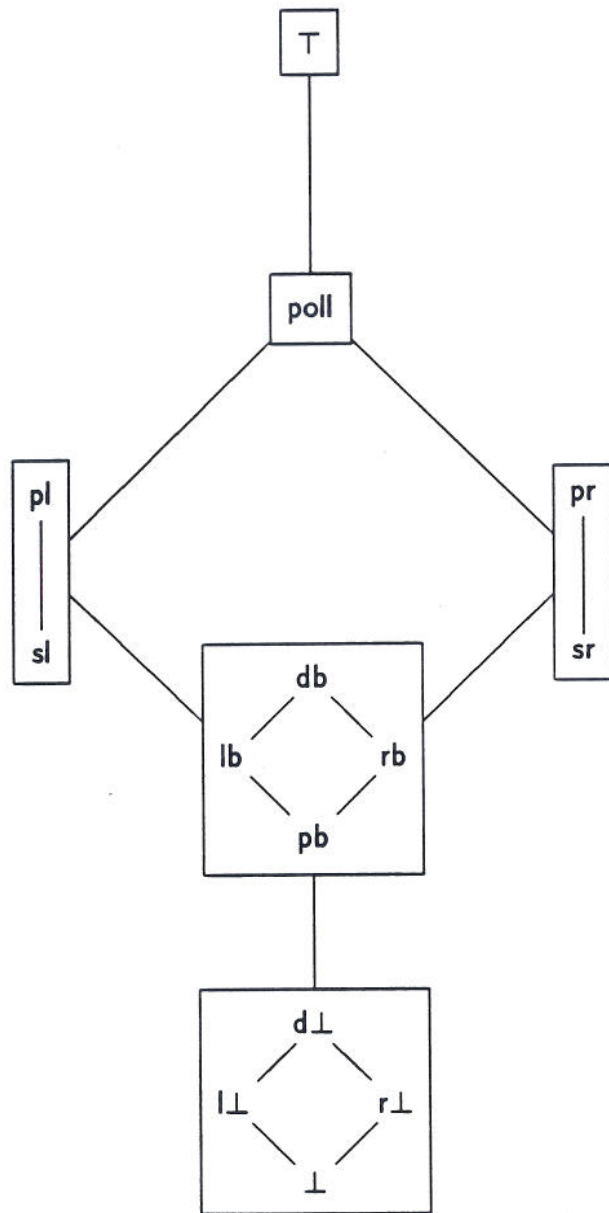


Figure 13. Quotient of Figure 8 by extensional equivalence.

$=^e_k$ be the equivalence relation induced by γ_k . For example, we have, for each $k \geq 0$:

$$\begin{aligned} \text{fun}_k(b_n) &= b && \text{if } k \leq n \\ &= \perp && \text{if } k > n. \end{aligned}$$

Clearly, $b_n =^e_k b_{n+1}$ iff $k \neq n + 1$.

Again the Kleisli category quotients onto the underlying category under the congruence induced by γ_k . Figure 14 shows the quotient of Figure 9 under the equivalence induced by γ_{n+1} .

5 Products and Exponentiation

5.1 Products

Now suppose that the underlying category \mathcal{C} has products, and for each pair of objects A_1 and A_2 there is a distinguished product, which we denote $A_1 \times A_2$, with π_i ($i = 1, 2$) being the projections. It is easy to show that distinguished product objects in \mathcal{C} are also product objects in \mathcal{C}_T , with projections in \mathcal{C}_T given by:

$$\begin{aligned} \hat{\pi}_i &: A_1 \times A_2 \rightarrow^{\mathcal{C}_T} A_i \\ \hat{\pi}_i &= \epsilon_{A_i} \circ T \pi_i \\ &= \pi_i \circ \epsilon_{A_1 \times A_2}. \end{aligned}$$

Pairing of morphisms in \mathcal{C}_T is the pairing of morphisms in \mathcal{C} , and the combination $\langle T \pi_1, T \pi_2 \rangle$ plays a special rôle in light of the following properties.

Proposition 5.1 *If \mathcal{C} has products then $\sigma : T(- \times -) \rightarrow T(-) \times T(-)$ defined by:*

$$\begin{aligned} \sigma_{A,B} &: T(A \times B) \rightarrow TA \times TB \\ \sigma_{A,B} &= \langle T \pi_1, T \pi_2 \rangle \end{aligned}$$

is a natural transformation such that, for all objects A and B ,

$$\begin{aligned} (\epsilon_A \times \epsilon_B) \circ \sigma_{A,B} &= \epsilon_{A \times B} \\ (\delta_A \times \delta_B) \circ \sigma_{A,B} &= \sigma_{TA, TB} \circ T \sigma_{A,B} \circ \delta_{A \times B} \\ \sigma_{A,B} \circ \gamma_{A \times B} &= \gamma_A \times \gamma_B. \end{aligned}$$

Proof: Naturality of σ follows from the universal property of products and the functoriality of T . The remaining properties are easy consequences of naturality of ϵ , δ , and γ . Note the identity $\sigma \circ T \langle f, g \rangle = \langle Tf, Tg \rangle$. In particular, $\sigma \circ T \sigma = \langle T^2 \pi_1, T^2 \pi_2 \rangle$. \square

BROOKES AND GEVA : COMPUTATIONAL COMONADS

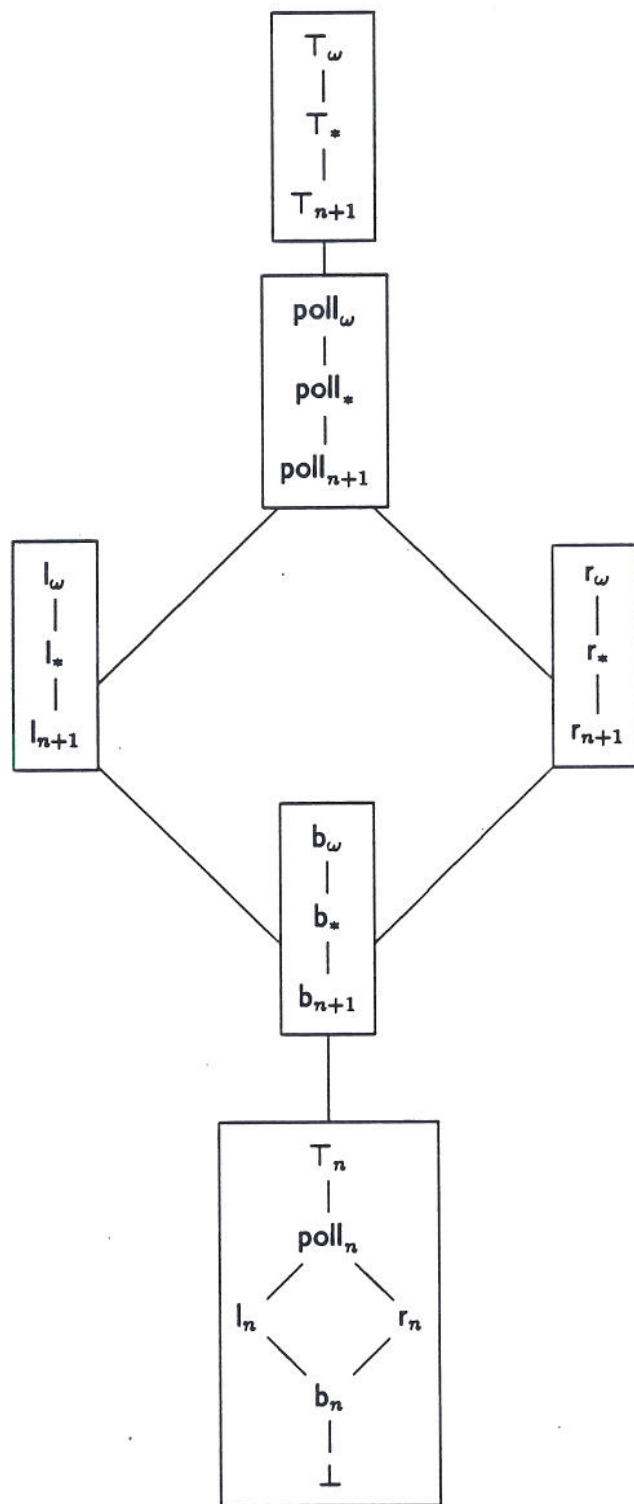


Figure 14. Quotient of Figure 9 by the equivalence induced by γ_{n+1} .

5.2 Exponentiation

Now suppose that the underlying category \mathcal{C} is cartesian closed. That is, we assume that \mathcal{C} has a distinguished terminal object and distinguished binary products, and that for every pair of objects B and C there is a distinguished exponentiation object $B \rightarrow C$ satisfying the usual requirements: for all B and C there is a morphism $\text{app}_{B,C} : (B \rightarrow C) \times B \rightarrow^{\mathcal{C}} C$ such that, for all A and all morphisms $f : A \times B \rightarrow^{\mathcal{C}} C$ there is a unique morphism $\text{curry}(f) : A \rightarrow^{\mathcal{C}} (B \rightarrow C)$ satisfying $\text{app}_{B,C} \circ (\text{curry}(f) \times \text{id}_B) = f$.

Equivalently, a category is cartesian closed if it has finite products and there is a pair of natural isomorphisms

$$\begin{aligned} \text{curry} : \mathcal{C}(- \times B, C) &\xrightarrow{\sim} \mathcal{C}(-, B \rightarrow C) \\ \text{uncurry} : \mathcal{C}(-, B \rightarrow C) &\xrightarrow{\sim} \mathcal{C}(- \times B, C). \end{aligned}$$

Here $\mathcal{C}(- \times B, C)$ and $\mathcal{C}(-, B \rightarrow C)$ are contravariant hom-functors from \mathcal{C}^{op} to the category *Set*, with the standard definitions [1]. This is the same as requiring that $\text{curry}(\text{uncurry } h) = h$ and $\text{uncurry}(\text{curry } g) = g$, together with the following naturality conditions: for all $f : A \rightarrow^{\mathcal{C}} A'$, $g : A' \times B \rightarrow^{\mathcal{C}} C$, and $h : A' \rightarrow^{\mathcal{C}} (B \rightarrow C)$,

$$\begin{aligned} \text{curry}(g \circ (f \times \text{id})) &= (\text{curry } g) \circ f \\ \text{uncurry}(h) \circ (f \times \text{id}) &= \text{uncurry}(h \circ f). \end{aligned}$$

It follows easily from these conditions that one can choose $\text{app} = \text{uncurry}(\text{id})$ as a suitable application morphism.

We want to give some general conditions under which analogous properties can be obtained for the Kleisli category \mathcal{C}_T . Assuming that \mathcal{C} is cartesian closed, the obvious candidate in \mathcal{C}_T for the exponential object of B and C is $TB \rightarrow C$. Moreover, we know that there is a natural isomorphism between $\mathcal{C}_T(A, TB \rightarrow C)$ and $\mathcal{C}(TA \times TB, C)$. Since $\mathcal{C}_T(A \times B, C)$ is $\mathcal{C}(T(A \times B), C)$, it is clear that we must make some assumptions about the relationship between $T(A \times B)$ and $TA \times TB$.

If $T(A \times B)$ and $TA \times TB$ are naturally isomorphic it is easy to show that \mathcal{C}_T is cartesian closed whenever \mathcal{C} is. This is apparently a ‘‘Folk Theorem’’. The comonad T_1 has this property, and we gave in [4] a proof using this property that the Kleisli category of T_1 is cartesian closed.

However, there are reasonable examples in which T does not preserve products, including T_2 and T_3 as described earlier. Instead, we will make a weaker assumption: that the comonad can be equipped with natural ways to move

back and forth between $T(A \times B)$ and $TA \times TB$ that interact sensibly with the comonad operations ϵ , δ , and γ . This can be conveniently summarized by means of two natural transformations *split* and *merge* satisfying certain combinational laws.

Definition 5.2 Let $(T, \epsilon, \delta, \gamma)$ be a computational comonad. A *computational pairing* is a pair of natural transformations

$$\begin{aligned} \text{split} &: T(- \times -) \rightarrow T(-) \times T(-) \\ \text{merge} &: T(-) \times T(-) \rightarrow T(- \times -) \end{aligned}$$

such that, for all objects A and B , the following properties hold:

$$\begin{aligned} (\epsilon_A \times \epsilon_B) \circ \text{split}_{A,B} &= \epsilon_{A \times B} \\ \epsilon_{A \times B} \circ \text{merge}_{A,B} &= \epsilon_A \times \epsilon_B \\ \text{split}_{A,B} \circ \gamma_{A \times B} &= \gamma_A \times \gamma_B \\ \text{merge}_{A,B} \circ (\gamma_A \times \gamma_B) &= \gamma_{A \times B} \\ (\delta_A \times \delta_B) \circ \text{split}_{A,B} &= \text{split}_{TA, TB} \circ T \text{split}_{A,B} \circ \delta_{A \times B} \\ \text{merge}_{TA, TB} \circ (\delta_A \times \delta_B) &= T \text{split}_{A,B} \circ \delta_{A \times B} \circ \text{merge}_{A,B}. \end{aligned}$$

Naturality of *split* and *merge* requires that for $f : A \rightarrow^c A'$ and $g : B \rightarrow^c B'$,

$$\begin{aligned} \text{split}_{A', B'} \circ T(f \times g) &= (Tf \times Tg) \circ \text{split}_{A,B} \\ \text{merge}_{A', B'} \circ (Tf \times Tg) &= T(f \times g) \circ \text{merge}_{A,B}. \end{aligned}$$

We summarize these properties in diagram form in Figure 15. •




The properties listed above formalize the sense in which we require the splitting and merging operations to interact sensibly with ϵ , δ , and γ . In particular, the following properties follow immediately.

Corollary 5.3 Let $(T, \epsilon, \delta, \gamma)$ be a computational comonad and let *split* and *merge* form a computational pairing. Then for all A and B ,

$$\begin{aligned} (\epsilon_A \times \epsilon_B) \circ \text{split}_{A,B} \circ \text{merge}_{A,B} &= (\epsilon_A \times \epsilon_B) \\ \text{split}_{A,B} \circ \text{merge}_{A,B} \circ (\gamma_A \times \gamma_B) &= (\gamma_A \times \gamma_B) \\ \epsilon_{A \times B} \circ \text{merge}_{A,B} \circ \text{split}_{A,B} &= \epsilon_{A \times B} \\ \text{merge}_{A,B} \circ \text{split}_{A,B} \circ \gamma_{A \times B} &= \gamma_{A \times B}. \end{aligned}$$

We have already seen that $\sigma = \langle T \pi_1, T \pi_2 \rangle$ qualifies as a suitable split operation (Proposition 5.1). Despite this fact, *split* (and *merge*) are not generally uniquely determined by the computational pairing laws and we wish to permit the use of comonads with “non-standard” choices of *split*. Moreover, naturality of *split* and *merge* does not by itself imply the computational pairing laws.

BROOKES AND GEVA : COMPUTATIONAL COMONADS



5.3 Examples

1. Return again to the increasing path comonad T_1 . The natural transformation $\sigma = \langle T_1 \pi_1, T_1 \pi_2 \rangle$ is given by: $\sigma(u) = \langle \lambda n. \pi_1(u_n), \lambda n. \pi_2(u_n) \rangle$. This is actually an isomorphism, with inverse given by $\text{merge}(\langle s, t \rangle) = \lambda n. \langle s_n, t_n \rangle$. Intuitively, each of these two operations works “in parallel” on the two components.

Both σ and merge are natural transformations, and they satisfy the computational pairing laws, which state that:

- Merging and splitting commute with componentwise application of functions.
 - Merging and splitting respect lubs of sequences.
 - Merging and splitting respect prefixes.
 - Merging two degenerate computations produces a degenerate computation, and splitting a degenerate computation produces a pair of degenerate computations.
2. There are other intuitively sensible ways to split and merge in the increasing paths comonad T_1 . We can define a form of (left-first) interleaving by:

$$\text{lmerge}(\langle s, t \rangle) = \lambda n. \langle s_{\lfloor n/2 \rfloor}, t_{\lfloor n/2 \rfloor} \rangle.$$

For example, this gives:

$$\text{lmerge}(s_0 s_1 s_2^\omega, t_0 t_1 t_2^\omega) = \langle s_0, t_0 \rangle \langle s_1, t_0 \rangle \langle s_1, t_1 \rangle \langle s_2, t_1 \rangle \langle s_2, t_2 \rangle^\omega.$$

To go with lmerge , we define a split that operates only on alternate steps of a computation:

$$\text{split}_2(u) = \langle \lambda n. \pi_1(u_{2n}), \lambda n. \pi_2(u_{2n}) \rangle.$$

We then obtain the identity $\text{split}_2 \circ \text{lmerge} = \text{id}$.

It is easy to verify that lmerge and split_2 are natural transformations, and that the computational pairing properties hold, making use of the equalities $\lfloor \min(i, 2j)/2 \rfloor = \min(\lfloor i/2 \rfloor, j)$ and $\lceil \min(i, 2j)/2 \rceil = \min(\lceil i/2 \rceil, j)$.

There is clearly also a right-first version of interleaving rmerge and this also interacts sensibly with split_2 as given above.

3. For the strictly increasing paths comonad, each of the split-merge combinations above adapts in the obvious way, modified so as to ensure that the result of splitting a strictly increasing sequence of pairs is a pair of strictly increasing sequences. Thus, for example,

$$\begin{aligned} \sigma(\langle \perp, \perp \rangle \langle \top, \perp \rangle \langle \top, \top \rangle^\omega) &= \langle \perp \top^\omega, \perp \top^\omega \rangle \\ \text{merge}(\perp \top^\omega, \perp \top^\omega) &= \langle \perp, \perp \rangle \langle \top, \top \rangle^\omega. \end{aligned}$$

In fact $T_2(A \times B)$ and $T_2A \times T_2B$ are not generally isomorphic, because a strictly increasing sequence of pairs does not necessarily increase strictly in *both* components at each stage. Nevertheless, σ and merge are still natural transformations satisfying the requirements listed above for a computational pairing, and we have the identity $\sigma \circ \text{merge} = \text{id}$.

The (appropriately adjusted) lmerge and split_2 operations also form a computational pairing, and $\text{split}_2 \circ \text{lmerge} = \text{id}$; similar properties hold for rmerge and split_2 .

5.4 Pairing, currying and uncurrying on algorithms

Using the split operation of a computational pairing provides a way to combine a pair of algorithms into an algorithm on pairs. If split is taken to be σ , this is the standard way to form the product of two morphisms in the Kleisli category. We can also use split to define intensional analogues to the contravariant hom-functors $\mathcal{C}(- \times B, C)$ and $\mathcal{C}(-, B \rightarrow C)$.

Definition 5.4 Let \mathcal{C} be a category with finite products, let $(T, \epsilon, \delta, \gamma)$ be a computational comonad, and let split be a natural transformation from $T(- \times -)$ to $T- \times T-$. For $f : A \rightarrow^{\mathcal{C}_T} A'$ and $g : B \rightarrow^{\mathcal{C}_T} B'$ we define $(f \hat{\times} g) : (A \times B) \rightarrow^{\mathcal{C}_T} (A' \times B')$ by $f \hat{\times} g = (f \times g) \circ \text{split}$. •

Proposition 5.5 Let \mathcal{C} be a category with finite products, let $(T, \epsilon, \delta, \gamma)$ be a computational comonad and let split and merge form a computational pairing. Then there is a functor $\hat{\times}$ from $\mathcal{C}_T \times \mathcal{C}_T$ to \mathcal{C}_T that maps a pair of objects (A, B) to $A \times B$ and maps a pair of morphisms (a, b) to $(a \hat{\times} b) = (a \times b) \circ \text{split}$.

Proof: To show that $\hat{\times}$ maps identity morphisms to identity morphisms:

$$(\hat{\text{id}}_A \hat{\times} \hat{\text{id}}_B) = (\epsilon_A \times \epsilon_B) \circ \text{split} = \epsilon_{A \times B} = \hat{\text{id}}_{A \times B}.$$

BROOKES AND GEVA : COMPUTATIONAL COMONADS

To show that $\hat{\times}$ preserves composition, let $a : A \rightarrow^{c_T} A'$, $b : B \rightarrow^{c_T} B'$, $a' : A' \rightarrow^{c_T} A''$, and $b' : B' \rightarrow^{c_T} B''$. Then:

$$\begin{aligned}
 (a' \hat{\times} b') \bar{o} (a \hat{\times} b) &= ((a' \times b') \circ \text{split}) \bar{o} ((a \times b) \circ \text{split}) \\
 &= (a' \times b') \circ \text{split} \circ T(a \times b) \circ T \text{split} \circ \delta \\
 &= (a' \times b') \circ (Ta \times Tb) \circ \text{split} \circ T \text{split} \circ \delta \\
 &= (a' \times b') \circ (Ta \times Tb) \circ (\delta \times \delta) \circ \text{split} \\
 &= ((a' \circ Ta \circ \delta) \times (b' \circ Tb \circ \delta)) \circ \text{split} \\
 &= ((a' \bar{o} a) \times (b' \bar{o} b)) \circ \text{split} \\
 &= (a' \bar{o} a) \hat{\times} (b' \bar{o} b).
 \end{aligned}$$

□

Definition 5.6 Let \mathcal{C} be a cartesian closed category, let $(T, \epsilon, \delta, \gamma)$ be a computational comonad, and let **split** and **merge** form a computational pairing. The contravariant functor $\mathcal{C}_T(- \hat{\times} B, C)$ from $\mathcal{C}_T^{\text{op}}$ to Set is defined as follows.

- The functor maps an object A to $\mathcal{C}_T(A \times B, C)$.
- The functor maps a morphism $f : A \rightarrow^{c_T} A'$ to the function $\lambda g. g \bar{o} (f \hat{\times} \text{id}_B)$ from $\mathcal{C}_T(A' \times B, C)$ to $\mathcal{C}_T(A \times B, C)$.

Similarly, the contravariant functor $\mathcal{C}_T(-, TB \rightarrow C)$ from $\mathcal{C}_T^{\text{op}}$ to Set , is defined by:

- On objects the functor maps A to $\mathcal{C}_T(A, TB \rightarrow C)$.
- On morphisms the functor maps $f : A \rightarrow^{c_T} A'$ to the function $\lambda h. h \bar{o} f$ from $\mathcal{C}_T(A', TB \rightarrow C)$ to $\mathcal{C}_T(A, TB \rightarrow C)$.

Proposition 5.7 Let $(T, \epsilon, \delta, \gamma)$ be a computational comonad and let **split** and **merge** be a computational pairing. Given $a : T(A \times B) \rightarrow^c C$ and $b : TA \rightarrow^c (TB \rightarrow C)$, define

$$\begin{aligned}
 \widehat{\text{curry}}(a) : TA \rightarrow^c (TB \rightarrow C) & & \widehat{\text{uncurry}}(b) : T(A \times B) \rightarrow^c C \\
 \widehat{\text{curry}}(a) = \text{curry}(a \circ \text{merge}) & & \widehat{\text{uncurry}}(b) = \text{uncurry}(b) \circ \text{split}.
 \end{aligned}$$

Then:

- $\widehat{\text{curry}}$ and $\widehat{\text{uncurry}}$ are natural transformations:

$$\begin{aligned}\widehat{\text{curry}} &: \mathcal{C}_T(- \hat{\times} B, C) \rightarrow \mathcal{C}_T(-, TB \rightarrow C) \\ \widehat{\text{uncurry}} &: \mathcal{C}_T(-, TB \rightarrow C) \rightarrow \mathcal{C}_T(- \hat{\times} B, C).\end{aligned}$$

- For all $a : A \times B \rightarrow^{\mathcal{C}_T} C$, $\widehat{\text{uncurry}}(\widehat{\text{curry}}(a)) =^e a$.
- For all $f : A \times B \rightarrow^{\mathcal{C}} C$, $\widehat{\text{uncurry}}(\widehat{\text{curry}}(\text{alg } f)) = \text{alg } f$.

Proof:

- Naturality of $\widehat{\text{curry}}$ requires that $\widehat{\text{curry}}(g \bar{\circ} (f \hat{\times} \widehat{\text{id}}_B)) = (\widehat{\text{curry}} g) \bar{\circ} f$, for all $f : A \rightarrow^{\mathcal{C}_T} A'$ and $g : A' \times B \rightarrow^{\mathcal{C}_T} C$. This follows from naturality of currying in \mathcal{C} and the properties of computational pairing:

$$\begin{aligned}\widehat{\text{curry}}(g \bar{\circ} (f \hat{\times} \widehat{\text{id}})) &= \widehat{\text{curry}}(g \circ T(f \times \epsilon) \circ T \text{ split} \circ \delta) \\ &= \text{curry}(g \circ T(f \times \epsilon) \circ T \text{ split} \circ \delta \circ \text{merge}) \\ &= \text{curry}(g \circ T(f \times \epsilon) \circ \text{merge} \circ (\delta \times \delta)) \\ &= \text{curry}(g \circ \text{merge} \circ (Tf \times T\epsilon) \circ (\delta \times \delta)) \\ &= \text{curry}(g \circ \text{merge} \circ ((Tf \circ \delta) \times (T\epsilon \circ \delta))) \\ &= \text{curry}(g \circ \text{merge} \circ ((Tf \circ \delta) \times \text{id})) \\ &= \text{curry}(g \circ \text{merge}) \circ (Tf \circ \delta) \\ &= \widehat{\text{curry}}(g) \circ Tf \circ \delta \\ &= \widehat{\text{curry}}(g) \bar{\circ} f.\end{aligned}$$

- Similarly, to show naturality of $\widehat{\text{uncurry}}$ we need $\widehat{\text{uncurry}}(h) \bar{\circ} (f \hat{\times} \widehat{\text{id}}) = \widehat{\text{uncurry}}(h \bar{\circ} f)$, for all $f : A \rightarrow^{\mathcal{C}_T} A'$ and $h : A' \rightarrow^{\mathcal{C}_T} (TB \rightarrow C)$. Again the proof is straightforward:

$$\begin{aligned}\widehat{\text{uncurry}}(h) \bar{\circ} (f \hat{\times} \widehat{\text{id}}) &= \text{uncurry}(h) \circ \text{split} \circ T(f \times \epsilon) \circ T \text{ split} \circ \delta \\ &= \text{uncurry}(h) \circ (Tf \times T\epsilon) \circ \text{split} \circ T \text{ split} \circ \delta \\ &= \text{uncurry}(h) \circ (Tf \times T\epsilon) \circ (\delta \times \delta) \circ \text{split} \\ &= \text{uncurry}(h) \circ ((Tf \circ \delta) \times (T\epsilon \circ \delta)) \circ \text{split} \\ &= \text{uncurry}(h) \circ ((Tf \circ \delta) \times \text{id}) \circ \text{split} \\ &= \text{uncurry}(h \circ Tf \circ \delta) \circ \text{split} \\ &= \text{uncurry}(h \bar{\circ} f) \circ \text{split} \\ &= \widehat{\text{uncurry}}(h \bar{\circ} f).\end{aligned}$$

- Let $a : A \times B \rightarrow^{\mathcal{C}_T} C$. Then

$$\begin{aligned}\widehat{\text{uncurry}}(\widehat{\text{curry}}(a)) \circ \gamma &= \text{uncurry}(\text{curry}(a \circ \text{merge})) \circ \text{split} \circ \gamma \\ &= (a \circ \text{merge}) \circ \text{split} \circ \gamma \\ &= a \circ (\text{merge} \circ \text{split} \circ \gamma) \\ &= a \circ \gamma,\end{aligned}$$

showing that $\widehat{\text{uncurry}}(\widehat{\text{curry}} a) =^e a$.

- Let $f : A \times B \rightarrow^C C$. Then

$$\begin{aligned}
 \widehat{\text{uncurry}}(\widehat{\text{curry}}(\text{alg } f)) &= \text{uncurry}(\text{curry}(\text{alg}(f) \circ \text{merge})) \circ \text{split} \\
 &= \text{alg}(f) \circ \text{merge} \circ \text{split} \\
 &= f \circ \epsilon \circ \text{merge} \circ \text{split} \quad \text{by Corollary 5.3} \\
 &= f \circ \epsilon \\
 &= \text{alg } f.
 \end{aligned}$$

□

We have shown that an intensional pairing produces a weak form of exponentiation structure: we obtain notions of currying and uncurrying on algorithms that are natural transformations but satisfy a weaker condition than isomorphism. We may rephrase these properties in terms of the existence of a notion of “application” in the intensional category as follows.

Proposition 5.8 *Let \mathcal{C} be a cartesian closed category, $(T, \epsilon, \delta, \gamma)$ be a computational comonad and let split and merge be a computational pairing. For all B and C there is an “application morphism”*

$$\widehat{\text{app}}_{B,C} : [TB \rightarrow C] \times B \rightarrow^{c_T} C$$

such that, for all $a : A \times B \rightarrow^{c_T} C$

$$\widehat{\text{app}}_{B,C} \circ \widehat{\text{curry}}(a) \hat{\times} \widehat{\text{id}}_B =^e a.$$

Proof: Define $\widehat{\text{app}}_{B,C} = \widehat{\text{uncurry}}(\widehat{\text{id}}_{TB \rightarrow C}) = \widehat{\text{uncurry}}(\epsilon_{TB \rightarrow C})$. As a corollary of the naturality of $\widehat{\text{uncurry}}$ (Proposition 5.7), we get:

$$\begin{aligned}
 \widehat{\text{app}} \circ \widehat{\text{curry}}(b \hat{\times} \widehat{\text{id}}) &= \widehat{\text{uncurry}}(\widehat{\text{id}}) \circ \widehat{\text{curry}}(b \hat{\times} \widehat{\text{id}}) \\
 &= \widehat{\text{uncurry}}(\widehat{\text{id}} \circ b) \\
 &= \widehat{\text{uncurry}}(b).
 \end{aligned}$$

Thus, in particular, $\widehat{\text{app}} \circ \widehat{\text{curry}}(a) \hat{\times} \widehat{\text{id}} = \widehat{\text{uncurry}}(\widehat{\text{curry}}(a)) =^e a$.

Note that although $\widehat{\text{curry}}(a)$ is not necessarily the only morphism h such that $\widehat{\text{app}} \circ \widehat{\text{curry}}(h) \hat{\times} \widehat{\text{id}} =^e a$, all such morphisms satisfy the condition that $\widehat{\text{uncurry}}(h) =^e a$. □

Thus, we have a weak form of cartesian closedness: instead of the usual diagram for exponentiation we replace $=$ by $=^e$ and we relax the uniqueness condition. This is summarized in Figure 16.

Next we consider what happens if we make further assumptions on the relationship between split and merge.

$$\begin{array}{ccc}
 A \times B & \xrightarrow{a} & C \\
 \downarrow \widehat{\text{curry}}(a) \hat{\times} \widehat{\text{id}} & \overset{=^e}{\nearrow} & \nearrow \widehat{\text{app}} \\
 (TB \rightarrow C) \times B & &
 \end{array}$$

Figure 16. When T has computational pairing this diagram in \mathcal{C}_T commutes up to extensional equivalence, for all $a : A \times B \rightarrow^{\mathcal{C}_T} C$.

Proposition 5.9 *Let \mathcal{C} be a cartesian closed category and $(T, \epsilon, \delta, \gamma)$ be a computational comonad with a computational pairing split and merge.*

- *If $\text{merge} \circ \text{split} = \text{id}$ then $\widehat{\text{uncurry}} \circ \widehat{\text{curry}} = \text{id}$.*
- *If $\text{split} \circ \text{merge} = \text{id}$ then $\widehat{\text{curry}} \circ \widehat{\text{uncurry}} = \text{id}$.*

As a corollary we get the following version of the ‘‘Folk Theorem’’:

Corollary 5.10 *If \mathcal{C} is cartesian closed and $(T, \epsilon, \delta, \gamma)$ is a computational comonad over \mathcal{C} , with a computational pairing such that $\text{merge} \circ \text{split} = \text{id}$ and $\text{split} \circ \text{merge} = \text{id}$, then the category \mathcal{C}_T is cartesian closed.*

Note the important fact that our definitions are parameterized by the choice of split and merge . Once these are chosen, $\widehat{\text{app}}$, $\widehat{\text{curry}}$ and $\widehat{\text{uncurry}}$ are determined uniquely. The Kleisli category itself is independent of split and merge ; what happens, however, is that each choice of these two natural transformations induces a (weak form of) exponentiation structure on this category. The Kleisli category may possess many different notions of merging and splitting, and therefore many different ways to curry, uncurry and apply algorithms. This means that one may use the Kleisli category to give an interpretation to a functional programming language containing several syntactically and semantically distinct forms of application. This would be desirable, for instance, if the language included both a strict and a non-strict form of application.

5.5 Examples

1. The Kleisli category based on the increasing path comonad T_1 is cartesian closed, with exponentiation structure built from the standard split-merge combination, which form an isomorphism.

Using the computational pairing lmerge and split_2 , we obtain intensional forms of currying, uncurrying, and application which we will call $\widehat{\text{curry}}_1$, $\widehat{\text{uncurry}}_2$ and $\widehat{\text{app}}_1$. This provides a weak form of exponentiation: $\widehat{\text{curry}}_1$ and $\widehat{\text{uncurry}}_2$ are natural transformations, and for all $a : A \times B \rightarrow^{c_{T_1}} C$ we get

$$\widehat{\text{app}}_1 \circ (\widehat{\text{curry}}_1(a) \hat{\times} \widehat{\text{id}}) =^e a.$$

Since $\text{split}_2 \circ \text{lmerge} = \text{id}$, it follows that $\widehat{\text{curry}}_1(\widehat{\text{uncurry}}_2 h) = h$. However, $\widehat{\text{uncurry}}_2(\widehat{\text{curry}}_1 g) =^e g$, and equality may fail here. For example,

$$\widehat{\text{uncurry}}_2(\widehat{\text{curry}}_1 \text{pb}) = \widehat{\text{uncurry}}_2(\widehat{\text{curry}}_1 \text{lb}) = \text{pb},$$

and $\widehat{\text{uncurry}}_2(\widehat{\text{curry}}_1 \text{rb})$ is the least algorithm mapping the path $\langle \perp, \top \rangle \langle \perp, \top \rangle \langle \top, \top \rangle^\omega$ to \top . This algorithm of course also computes the function b .

Similar properties hold for the computation pairing rmerge and split_2 , with the derived operations $\widehat{\text{curry}}_r$, $\widehat{\text{uncurry}}_2$ and $\widehat{\text{app}}_r$.

2. The strictly increasing paths comonad T_2 , with computational pairing σ and merge , again has operations $\widehat{\text{curry}}$, $\widehat{\text{uncurry}}$ and $\widehat{\text{app}}$ that provide weak forms of exponentiation. As we remarked earlier, σ and merge are not isomorphisms. Instead, $\sigma \circ \text{merge} = \text{id}$ and for all $u \in T_2(A \times B)$ the computation $\text{merge}(\sigma(u))$ is pointwise above u . Hence, $\widehat{\text{curry}}(\widehat{\text{uncurry}} h) = h$ and $\widehat{\text{uncurry}}(\widehat{\text{curry}} g) =^e g$. As an example, we have:

$$\widehat{\text{uncurry}}(\widehat{\text{curry}} \text{lb}) = \widehat{\text{uncurry}}(\widehat{\text{curry}} \text{rb}) = \widehat{\text{uncurry}}(\widehat{\text{curry}} \text{pb}) = \text{pb}.$$

The lmerge and split_2 computational pairing also gives rise to a weak form of exponentiation, as does the rmerge and split_2 computational pairing.

6 Ordered categories

So far, although our principal examples were based on a cartesian closed category of domains, we have not fully exploited the order structure. This permitted us to state and prove results that hold in a more general category-theoretic setting. Next we suppose that the underlying category is an ordered category: each hom-set is equipped with a complete partial order, and composition is continuous. A functor T of ordered categories is required to respect the ordering, in that for all $f, g : A \rightarrow^c B$ if $f \leq g$ then $Tf \leq Tg$. Moreover, T must also be continuous (in its action on morphisms). All of our examples so far satisfy these conditions.

BROOKES AND GEVA : COMPUTATIONAL COMONADS

Suppose $(T, \epsilon, \delta, \gamma)$ is a computational comonad over an ordered category \mathcal{C} . Then clearly \mathcal{C}_T is again an ordered category. All of the results of the previous sections go through in the ordered setting. In particular, the functors fun and alg introduced earlier respect the ordering; and the proof of Proposition 4.2 can be adapted to show that the extensional quotient of the ordering on $\mathcal{C}_T(A, B)$ is just the order on $\mathcal{C}(A, B)$.

We can also obtain some slightly stronger results by taking advantage of the ordering. We omit most of the proofs, which may be easily obtained from the results above, using monotonicity and continuity of composition.

Proposition 6.1 *Let \mathcal{C} be an ordered ccc and let $(T, \epsilon, \delta, \gamma)$ be a computational comonad over \mathcal{C} with a computational pairing.*

- *If $\text{split} \circ \text{merge} \geq \text{id}$ then $\widehat{\text{curry}} \circ \widehat{\text{uncurry}} \geq \text{id}$.*
- *If $\text{split} \circ \text{merge} \leq \text{id}$ then $\widehat{\text{curry}} \circ \widehat{\text{uncurry}} \leq \text{id}$.*
- *If $\text{merge} \circ \text{split} \geq \text{id}$ then $\widehat{\text{uncurry}} \circ \widehat{\text{curry}} \geq \text{id}$.*
- *If $\text{merge} \circ \text{split} \leq \text{id}$ then $\widehat{\text{uncurry}} \circ \widehat{\text{curry}} \leq \text{id}$.*

Next we introduce a simple generalization of the notion of cartesian closed ordered category, obtained by relaxing the requirement that currying and uncurrying form an isomorphism. Instead we allow currying and uncurrying to form an adjunction in each homset, so that we have an example of a *local adjunction* (see for example [9]) with additional properties. The relevance of “lax” notions of adjunction such as these in computational settings (albeit with different motivations) has been pointed out in different contexts by other authors, for instance in [14].

Definition 6.2 An ordered category \mathcal{C} is *cartesian up-closed* if and only if it has finite products and for all pairs of objects B and C there is an object $B \rightarrow C$ and a pair of lax natural transformations curry , uncurry between $\mathcal{C}(- \times B, C)$ and $\mathcal{C}(-, B \rightarrow C)$ satisfying:

$$\begin{aligned}
 \text{curry}(\text{uncurry } h) &\leq h \\
 \text{uncurry}(\text{curry } g) &\geq g \\
 \text{curry}(g \circ (f \times \text{id})) &\leq \text{curry}(g) \circ f \\
 \text{uncurry}(h) \circ (f \times \text{id}) &\geq \text{uncurry}(h \circ f).
 \end{aligned}$$

Similarly, we say that \mathcal{C} is *cartesian down-closed* iff it has finite products and there is a pair of lax natural transformations curry , uncurry satisfying:

$$\begin{aligned} \text{curry}(\text{uncurry } h) &\geq h \\ \text{uncurry}(\text{curry } g) &\leq g \\ \text{curry}(g \circ (f \times \text{id})) &\geq \text{curry}(g) \circ f \\ \text{uncurry}(h) \circ (f \times \text{id}) &\leq \text{uncurry}(h \circ f). \end{aligned}$$

Definition 6.3 Let \mathcal{C} be an ordered category with finite products, and let B and C be objects of \mathcal{C} .

- An *up-exponential* for B and C is an object $B \rightarrow C$ of \mathcal{C} together with a morphism $\text{app}_{B,C} : (B \rightarrow C) \times B \rightarrow^{\mathcal{C}} C$ such that for every $f : A \times B \rightarrow^{\mathcal{C}} C$ there is a least morphism $\text{curry}(f) : A \rightarrow^{\mathcal{C}} (B \rightarrow C)$ satisfying:

$$\text{app} \circ (\text{curry}(f) \times \text{id}) \geq f.$$

- A *down-exponential* for B and C is an object $B \rightarrow C$ of \mathcal{C} together with a morphism $\text{app}_{B,C} : (B \rightarrow C) \times B \rightarrow^{\mathcal{C}} C$ such that for every $f : A \times B \rightarrow^{\mathcal{C}} C$ there is a greatest morphism $\text{curry}(f) : A \rightarrow^{\mathcal{C}} (B \rightarrow C)$ satisfying:

$$\text{app} \circ (\text{curry}(f) \times \text{id}) \leq f.$$

The following result may be shown by adapting the usual proof that the two alternative definitions of cartesian closed categories are equivalent.

Proposition 6.4 *An ordered category \mathcal{C} is cartesian up-closed iff it has finite products and up-exponentials.*

An ordered category is cartesian down-closed iff it has finite products and down-exponentials.

Note that in general up-exponentials and down-exponentials need not be unique (even up to isomorphism), since their definition relies on the choice of computational pairing. However, if the same object $B \rightarrow C$ and morphism $\text{app}_{B,C}$ qualifies simultaneously as an up- and a down-exponential then it forms the usual notion of exponentiation and the category is cartesian closed in the usual sense.

Proposition 6.5 *Let \mathcal{C} be a cartesian up-closed category and $(T, \epsilon, \delta, \gamma)$ be a computational comonad with a computational pairing such that*

$$\begin{aligned} \text{split} \circ \text{merge} &\leq \text{id} \\ \text{merge} \circ \text{split} &\geq \text{id} \\ (\delta \times \delta) \circ \text{split} &\leq \text{split} \circ T \text{ split} \circ \delta \\ \text{merge} \circ (\delta \times \delta) &\geq T \text{ split} \circ \delta \circ \text{merge}. \end{aligned}$$

Then \mathcal{C}_T is cartesian up-closed.

Proof:

- Let $a : A \times B \rightarrow^{\mathcal{C}_T} C$ and $b : A \rightarrow^{\mathcal{C}_T} (TB \rightarrow C)$. Then:

$$\begin{aligned} \widehat{\text{uncurry}}(\widehat{\text{curry}} a) &= \text{uncurry}(\text{curry}(a \circ \text{merge})) \circ \text{split} \\ &\geq a \circ \text{merge} \circ \text{split} \geq a \\ \widehat{\text{curry}}(\widehat{\text{uncurry}} b) &= \text{curry}(\text{uncurry}(b) \circ \text{split} \circ \text{merge}) \\ &\leq \text{curry}(\text{uncurry} b) \leq b. \end{aligned}$$

- To show that $\widehat{\text{curry}}$ is a lax natural transformation, let $f : A \rightarrow^{\mathcal{C}_T} A'$ and $g : A' \times B \rightarrow^{\mathcal{C}_T} C$. Then:

$$\begin{aligned} \widehat{\text{curry}}(g \circ (f \hat{\times} \text{id})) &= \widehat{\text{curry}}(g \circ T(f \times \epsilon) \circ T \text{ split} \circ \delta) \\ &= \text{curry}(g \circ T(f \times \epsilon) \circ T \text{ split} \circ \delta \circ \text{merge}) \\ &\leq \text{curry}(g \circ T(f \times \epsilon) \circ \text{merge} \circ (\delta \times \delta)) \\ &= \text{curry}(g \circ \text{merge} \circ (Tf \times T\epsilon) \circ (\delta \times \delta)) \\ &= \text{curry}(g \circ \text{merge} \circ ((Tf \circ \delta) \times (T\epsilon \circ \delta))) \\ &= \text{curry}(g \circ \text{merge} \circ ((Tf \circ \delta) \times \text{id})) \\ &\leq \text{curry}(g \circ \text{merge}) \circ (Tf \circ \delta) \\ &= \widehat{\text{curry}}(g) \circ Tf \circ \delta \\ &= \widehat{\text{curry}}(g) \circ f. \end{aligned}$$

- To show that $\widehat{\text{uncurry}}$ is a lax natural transformation, suppose that $f : A \rightarrow^{\mathcal{C}_T} A'$ and $h : A' \rightarrow^{\mathcal{C}_T} (TB \rightarrow C)$. Then

$$\begin{aligned} \widehat{\text{uncurry}}(h) \circ (f \hat{\times} \text{id}) &= \text{uncurry}(h) \circ \text{split} \circ T(f \times \epsilon) \circ T \text{ split} \circ \delta \\ &\geq \text{uncurry}(h) \circ (Tf \times T\epsilon) \circ \text{split} \circ T \text{ split} \circ \delta \\ &= \text{uncurry}(h) \circ (Tf \times T\epsilon) \circ (\delta \times \delta) \circ \text{split} \\ &= \text{uncurry}(h) \circ ((Tf \circ \delta) \times (T\epsilon \circ \delta)) \circ \text{split} \\ &= \text{uncurry}(h) \circ ((Tf \circ \delta) \times \text{id}) \circ \text{split} \\ &\geq \text{uncurry}(h \circ Tf \circ \delta) \circ \text{split} \\ &= \text{uncurry}(h \circ f) \circ \text{split} \\ &= \widehat{\text{uncurry}}(h \circ f). \end{aligned}$$

□

A similar result holds for a cartesian down-closed category with a computational pairing satisfying reversed inequalities.

7 Examples

We now return to the third comonad introduced earlier, after which we will introduce briefly some related notions of computation on different categories of domains and functions.

7.1 Timed data

In the timed data comonad T_3 , the standard split operation is:

$$\text{split} \langle \langle a, b \rangle, n \rangle = \langle \langle a, n \rangle, \langle b, n \rangle \rangle.$$

Given our interpretation of $\langle d, n \rangle$ as a computation yielding d at cost n , an obvious choice for a merge operation is:

$$\text{merge}(\langle a, m \rangle, \langle b, n \rangle) = \langle \langle a, b \rangle, \max(m, n) \rangle.$$

Both of these operations are natural transformations, and we obtain the following properties:

$$\begin{aligned} \text{split} \circ \text{merge} &\sqsubseteq \text{id} \\ \text{merge} \circ \text{split} &= \text{id} \\ (\delta \times \delta) \circ \text{split} &= \text{split} \circ T_3 \text{ split} \circ \delta \\ \text{merge} \circ (\delta \times \delta) &\sqsupseteq T_3 \text{ split} \circ \delta \circ \text{merge}. \end{aligned}$$

The underlying category is cartesian closed, hence also cartesian up-closed. It follows from Proposition 6.5 that the Kleisli category of T_3 is cartesian up-closed.

7.2 Strict algorithms

The category of domains and strict continuous functions is not cartesian closed, although the category does have products. For each pair of domains D and D' , the set of strict continuous functions $D \rightarrow D'$, ordered pointwise, is again a domain. The usual uncurrying operation on functions preserves

strictness, but the usual currying does not. Instead, we may define a variant form of currying by:

$$\begin{aligned} \text{curry}_s &: (A \times B \rightarrow_s C) \rightarrow (A \rightarrow_s (B \rightarrow_s C)) \\ \text{curry}_s(f) &= \lambda x. \lambda y. (x = \perp \vee y = \perp \rightarrow \perp, f(x, y)). \end{aligned}$$

When f is strict, $\text{curry}_s(f)$ is the best strict function approximating $\text{curry}(f)$ pointwise. For instance, let lor , ror and sor be the left-strict, right-strict, and doubly-strict or-functions. Then

$$\text{uncurry}(\text{curry}_s \text{lor}) = \text{uncurry}(\text{curry}_s \text{ror}) = \text{uncurry}(\text{curry}_s \text{sor}) = \text{sor}.$$

It is easy to check that curry_s is a natural transformation (and so is uncurry). For all $f : A \times B \rightarrow_s C$ and all $g : A \rightarrow_s (B \rightarrow_s C)$, we have:

$$\begin{aligned} \text{curry}_s(\text{uncurry } g) &= g \\ \text{uncurry}(\text{curry}_s f) &\sqsubseteq f. \end{aligned}$$

Hence, the category of domains and strict continuous functions is cartesian down-closed.

Let $T_1 D$ be the set of increasing paths over D (not just the strict continuous maps from $\mathbb{V}\text{Nat}$ to D), ordered pointwise. The maps ϵ , δ and γ are all strict, as are all of the split and merge operations above. We may therefore use the Kleisli construction to build a model of strict parallel algorithms. To illustrate this model, note that all of the algorithms of Figure 6 also belong in this category, with the exception of \top , which is non-strict.

Since the underlying category is cartesian down-closed, each of the computational pairings discussed earlier for T_1 gives rise to a down-exponentiation structure, so that the category of domains and strict algorithms is again cartesian down-closed.

We may also adapt the T_2 and T_3 comonads to this category.

7.3 Computation on effectively given domains

The category of effectively given domains and computable functions is cartesian closed. A domain is effectively given iff its finite elements are recursively enumerable (hence, countable), it is decidable whether two finite elements are consistent, and (an index for) the lub of two consistent finite elements is decidable (as a function of their indices). An element of D is computable iff the set of (indices of) its finite approximations is recursively enumerable.

The functor T_1 can be adapted to this category, by defining T_1D for an effectively given domain D to be the computable increasing paths over D (equivalently, the computable continuous functions from $\mathbb{V}\text{Nat}$ to D , ordered pointwise). All of the auxiliary operations (ϵ , δ , γ , and so on) are computable. Hence we obtain a category of effectively given domains and computable algorithms, and this category quotients onto the underlying category of effectively given domains and computable functions. This algorithms category is again cartesian closed.

The functor T_3 maps effectively given domains to effectively given domains, and again the auxiliary operations are computable. We therefore obtain a category of effectively given domains and T_3 -algorithms that quotients onto the underlying category and is cartesian up-closed.

The functor T_2 preserves algebraicity but not ω -algebraicity, since T_2D may have uncountably many finite elements. The T_2 comonad therefore does not adapt to the category of effectively given domains and computable functions.

7.4 Computation on pre-domains

We use the term *pre-domain* for a “bottomless” domain: a directed-complete, bounded complete, algebraic partial order with no requirement that there be a least element. The category of pre-domains and continuous functions is cartesian closed.

Let T_4D be the set of non-empty finite or infinite increasing sequences over D , ordered by the usual prefix ordering. Clearly this forms a pre-domain, and the finite elements are just the finite sequences. T_4D is generally a pre-domain rather than a domain, even if D has a least element, because the prefix ordering does not relate sequences with different first elements. We make T into a functor by specifying that (as usual) T_4f applies f componentwise.

Again we let ϵ be the lub operation and let δt be the sequence of (non-empty) prefixes of t . Then (T_4, ϵ, δ) forms a comonad.

We may regard a computation of length 1 as degenerate, and this corresponds to defining the function γ from D to T_4D by $\gamma d = \langle d \rangle$. Although this function γ is not continuous, so that we cannot claim that T_4 is a computational comonad, we still obtain a congruence relation on algorithms by defining

$$a =^e a' \iff \forall d \in D. a \langle d \rangle = a' \langle d \rangle.$$

Note that for all $f, g : A \rightarrow B$ we have

$$(f \circ \epsilon) =^e (g \circ \epsilon) \Rightarrow f = g.$$

It is then easy to modify the proof of Proposition 4.2 to show that the Kleisli category of this comonad quotients onto the underlying category under $=^e$.

We may define splitting and merging operations as follows. The standard way to split is:

$$\begin{aligned} \text{split}(\langle x_0, y_0 \rangle \dots \langle x_k, y_k \rangle) &= \langle x_0 \dots x_k, y_0 \dots y_k \rangle \\ \text{split}(\langle x_n, y_n \rangle_{n=0}^\infty) &= \langle \langle x_n \rangle_{n=0}^\infty, \langle y_n \rangle_{n=0}^\infty \rangle. \end{aligned}$$

Let merge be the least continuous function satisfying:

$$\begin{aligned} \text{merge}(x_0 \dots x_k, y_0 \dots y_m) &= \langle x_0, y_0 \rangle \dots \langle x_n, y_n \rangle & (n = \min(m, k)) \\ \text{merge}(x_0 \dots x_k, \langle y_i \rangle_{i=0}^\infty) &= \langle x_0, y_0 \rangle \dots \langle x_k, y_k \rangle \\ \text{merge}(\langle x_i \rangle_{i=0}^\infty, y_0 \dots y_m) &= \langle x_0, y_0 \rangle \dots \langle x_m, y_m \rangle \\ \text{merge}(\langle x_n \rangle_{n=0}^\infty, \langle y_n \rangle_{n=0}^\infty) &= \langle x_n, y_n \rangle_{n=0}^\infty. \end{aligned}$$

Clearly, $\text{merge} \circ \text{split} = \text{id}$ and $\text{split} \circ \text{merge} \sqsubseteq \text{id}$. These two operations are obviously natural and satisfy the computational pairing properties, except that $\epsilon \circ \text{merge} \sqsubseteq \epsilon \times \epsilon$. The Kleisli category is cartesian down-closed.

8 Conclusions

We have described a category-theoretic approach to intensional semantics, based on the idea that a notion of computation or intensional behavior may be modelled by means of a computational comonad, and that the Kleisli category thus obtained can be viewed as an intensional model. The morphisms in this category map computations to values, and from such a morphism one may recover a map from values to values. One may define an equivalence relation that identifies all algorithms that compute the same function, and this equivalence relation can be used to collapse the Kleisli category onto the underlying category.

We have identified a set of conditions under which the Kleisli category possesses exponentiations or weaker types of exponentiation, based on the existence of natural ways to pair computations. We described a series of examples to illustrate the applicability of our definitions and results. In doing so, we have placed our recent work [4] in a wider context.

Our work arose out of an attempt, begun in [3], to generalize an earlier intensional model of Berry and Curien [6]. They defined a category of deterministic concrete data structures and sequential algorithms, showed that this category is cartesian closed, and that it collapses onto the category of deterministic concrete data structures and sequential functions under an obvious notion of extensional equivalence. The sequential functions category

is not cartesian closed, and their construction of sequential algorithms was not based on a comonad. The operational semantics implicit in their work was coroutine-like, demand-driven, and lazy, but with the restriction that computation should proceed *sequentially*, with at most one argument being evaluated at a time. In our generalization of their model, we relax the sequentiality restriction so as to permit parallel computation, and we adopt an operational semantics based on parallel coroutine-like lazy evaluation. The *query model* of parallel algorithms between deterministic concrete data structures, described in [3], contains algorithms for non-sequential functions such as parallel-or. However, the model's construction was rather complex and we were unable to formulate a suitable notion of composition for algorithms. Instead, in [4] we presented a much more streamlined form of algorithm between Scott domains and for the first time we cast our construction in terms of a comonad. Of the comonads introduced in this paper, T_1 corresponds to the comonad used in [4]; T_2 is closer in spirit to the query model of [3], but we are able here to go considerably further. In [5] we show that the Berry-Curien sequential algorithms may be embedded in the T_1 model in a manner that respects operational behavior: the image of every sequential algorithm is a parallel algorithm that operates sequentially. Our model can thus be used to provide a semantics for the Berry-Curien language CDS0 of sequential algorithms. An extension of this to a parallel algorithmic language should be fairly straightforward.

Moggi has developed an abstract view of programming languages in which a notion of computation is modelled as a monad [12, 13]. Examples of notions of computation as monads include: computation with side-effects, computation with exceptions, partial computations, and non-deterministic computations. In this view, the meaning of a program is taken as a function from values to computations, and an intuitive operational semantics is that a program from A to B takes an input of type A and returns an output computation. This point of view is consistent with an *input-driven* lazy operational semantics. Our "opposite" point of view based on comonads (which are, after all, monads on the opposite category) is consistent with a *demand-driven* lazy operational semantics, as discussed above. Moggi states in [12] that his view of programs corresponds to call-by-value parameter passing, and he says that there is an alternative view of "programs as functions from computations to computations" corresponding to call-by-name. Our work shows that there is also a third alternative: programs as functions from computations to values.

Common to these approaches is the realization that values should be distinguished from computations. Both approaches use an endofunctor T to embody this distinction, TA representing a datatype of computations over A . Moggi shows that many such endofunctors arising in denotational semantics

can be usefully equipped with the extra structure needed to form a monad, and we show that certain endofunctors embodying intensional semantics can be usefully equipped with the structure of a comonad. The motivations and the operational intuitions behind the monad approach and the comonad approach are different, and the two approaches should be regarded as orthogonal or complementary. We plan to explore to what extent (and to what effect) the two approaches can be combined. For instance, given a comonad T and a monad P over the same category \mathcal{C} one might obtain (assuming that T and P satisfy certain properties) a category of (T, P) -algorithms, in which a morphism from A to B is a morphism in \mathcal{C} from TA to PB .

We plan to investigate notions of computation in further domain-theoretic settings. As noted above, we are already working on categories of algorithms on (generalized) concrete data structures [5]. It would be interesting to see if the Berry-Curien sequential algorithms category could be embedded in the Kleisli category of a suitable comonad over a sequential functions category, but no cartesian closed category based on sequential functions is yet known. We intend to investigate notions of computation on the category of dI-domains and stable functions [2], and on the category of qualitative domains and linear functions [7].

Acknowledgements

The diagrams in this paper were drawn using macros designed by John Reynolds.

References

- [1] M. Barr and C. Wells. *Category Theory for Computing Science*. Prentice-Hall International, 1990.
- [2] G. Berry. Stable models of typed λ -calculi. In *Proc. 5th Coll. on Automata, Languages and Programming*, number 62 in Lecture Notes in Computer Science, pages 72–89. Springer-Verlag, July 1978.
- [3] S. Brookes and S. Geva. Towards a theory of parallel algorithms on concrete data structures. In *Semantics for Concurrency, Leicester 1990*, pages 116–136. Springer Verlag, 1990.
- [4] S. Brookes and S. Geva. A cartesian closed category of parallel algorithms between Scott domains. Technical Report CMU-CS-91-159, Carnegie Mellon University, School of Computer Science, 1991. Submitted for publication.

BROOKES AND GEVA : COMPUTATIONAL COMONADS

- [5] S. Brookes and S. Geva. Continuous functions and parallel algorithms on concrete data structures. Technical Report CMU-CS-91-160, Carnegie Mellon University, School of Computer Science, 1991. To appear in the Proceedings of MFPS'91.
- [6] P.-L. Curien. *Categorical Combinators, Sequential Algorithms and Functional Programming*. Research Notes in Theoretical Computer Science. Pitman, 1986.
- [7] J.-Y. Girard. *Proofs and Types*. Cambridge University Press, 1989.
- [8] C. Gunter and D. Scott. Semantic domains. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*. MIT Press/Elsevier, 1990.
- [9] C. B. Jay. Local adjunctions. *Journal of Pure and Applied Logic*, 53:227-238, 1988.
- [10] G. Kahn and D. B. MacQueen. Coroutines and networks of parallel processes. In *Information Processing 1977*, pages 993-998. North Holland, 1977.
- [11] S. Mac Lane. *Categories for the Working Mathematician*. Springer-Verlag, 1971.
- [12] E. Moggi. Computational lambda-calculus and monads. In *Fourth Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1989.
- [13] E. Moggi. Notions of computation and monads. *Information and Computation*, 1991.
- [14] R. A. G. Seely. Modeling computations: A 2-categorical framework. In *Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1987.

